

(19) World Intellectual Property Organization
International Bureau



(43) International Publication Date
24 December 2003 (24.12.2003)

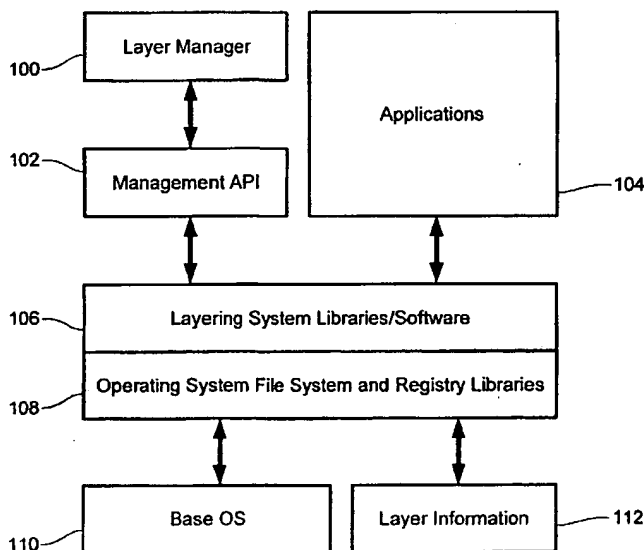
PCT

(10) International Publication Number
WO 03/107221 A1

- (51) International Patent Classification⁷: **G06F 17/30** (74) Agent: **ROBINSON, Everett, D.**; Parsons Behle & Latimer, 201 South Main Street, Suite 1800, Post Office Box 45898, Salt Lake City, UT 84145-0898 (US).
- (21) International Application Number: PCT/US03/18633
- (22) International Filing Date: 11 June 2003 (11.06.2003) (81) Designated States (*national*): AE, AG, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BY, BZ, CA, CH, CN, CO, CR, CU, CZ, DE, DK, DM, DZ, EC, EE, ES, FI, GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MA, MD, MG, MK, MN, MW, MX, MZ, NO, NZ, OM, PH, PL, PT, RO, RU, SC, SD, SE, SG, SK, SL, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, YU, ZA, ZM, ZW.
- (25) Filing Language: English
- (26) Publication Language: English
- (30) Priority Data:
60/387,969 12 June 2002 (12.06.2002) US
- (71) Applicant (*for all designated States except US*): **FS-LOGIC, INC.** [US/US]; 363 North University Ave., Suite 112, Provo, UT 84601 (US).
- (72) Inventors; and
- (75) Inventors/Applicants (*for US only*): **BLASER, Jared, R.** [US/US]; 419 North Palisade Drive, Orem, UT 84097 (US). **COOK, Randall, R.** [US/US]; 894 South 1650 East, Springville, UT 84663 (US).
- (84) Designated States (*regional*): ARIPO patent (GH, GM, KE, LS, MW, MZ, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HU, IE, IT, LU, MC, NL, PT, RO, SE, SI, SK, TR), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).
- Published:
— with international search report

[Continued on next page]

(54) Title: LAYERED COMPUTING SYSTEMS AND METHODS



(57) Abstract: The inventions relate generally to computer systems having facilities for providing virtual portions of file systems and configuration settings to applications. More particularly, the inventions relate to computer systems that provide a layer organization for files and configuration settings that can be overlaid on top of an operating system. An exemplary system includes a base operating system (110), operating system file system and registry libraries (108), layering system libraries and/or software (106), applications (104), a management applications programmer interface (102), and a layer manager application (100). Detailed information on various example embodiments of the inventions are provided in the Detailed Description below, and the inventions are defined by the appended claims.



— before the expiration of the time limit for amending the claims and to be republished in the event of receipt of amendments

For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.

[0001] Layered computing systems and methods

CROSS-REFERENCE TO RELATED APPLICATIONS

[0002] This application claims the benefit of U.S. Provisional Application No. 60/387,969 filed June 12, 2002 which is hereby incorporated by reference in its entirety.

BACKGROUND OF THE INVENTIONS

[0003] Prior computing systems have been susceptible to application conflicts with the host operating system (OS) and other applications. When an application is installed to an OS, a number of globally accessible files are often placed to the computing system, including for example shared libraries and system configuration. Those shared libraries are often provided in different versions, with applications requiring one version or another. A mismatch between a library version and a version required by an application sometimes results in that application crashing, becoming inoperable, or exhibiting other errors. Shared configuration elements are sometimes globally available to applications, which may write a favored configuration thereto. Following a write to that configuration other applications may be unable to read the configuration properly, or may be unable to function under a new specified configuration. Thus it is that following the installation of an application to a computer, other applications may stop working.

[0004] Installing a number of applications to a computer can be something of a black art. An administrator may, with good intentions and understanding, install several applications to a computer. Upon testing an installation or during use, the administrator or a user may discover that one or more applications operate errantly or not at all. It is usually not apparent which applications are in conflict. The administrator may enter a procedure in which applications are uninstalled from the computer in a process of elimination to find the offending applications. Sometimes de-installation programs do not remove all installed files, in which that procedure may fail to locate the problem. The administrator is then required to continue by creating a clean (or virgin) installation, and installing applications one at a time until the problem is located.

When applications are found to conflict, a choice must usually be made as to which one will be installed. One of the applications is sometimes installed to a different computer to avoid the conflict. If conflicting applications must be installed to a single computer, a new version of at least one of the applications must be sought and purchased from the software vendors. A non-

conflicting version may not be available, especially if a vendor is small, not supporting the application, or no longer in business.

[0005] Snapshot utilities are available, which generally operate to create a database of all files and registry settings on a computer. Prior to installing an application, a snapshot is taken of the files and registry settings. The application is then installed, and tested. If the application fails to work satisfactorily, the system can be restored by comparing the existing files and registry settings against the snapshot and removing installed files and otherwise restoring the system as before. Snapshot utilities have several limitations. First, if a newly installed application causes a prior installed application to fail, it is often not possible to simply revert to a snapshot made prior to older application installation, especially if there have been other applications installed in the interim. The administrator may be required to revert back to the earlier snapshot, and then re-install the intervening applications and the new application. Additionally, there are usually a limited number of snapshots that can be stored, and thus a required snapshot may not have been retained when found to be needed.

[0006] Likewise, a system may be restored to an earlier state if backups have been made. That restoration process, however, usually involves a significant amount of time and destroys all data recorded to the system after the time of the backup.

[0007] Another method involves recording a series of changes (or "diffs") to a buffer. Using that method a system can be restored back to a point in time by reverse application of the diffs to the file system back to the selected point in time. That method typically requires a fixed amount of disk space for the diff buffer, which becomes unavailable for regular use. As the buffer becomes full, the only way to continue to record diffs is to overwrite older diffs. Because of this limitation, the method can only restore a system back to a date for which diffs remain available. In addition, this method requires three disk operations per write request: one to read the existing disk information, one to write the diff, and one to write the original request. This method is therefore processor and disk intensive.

[0008] The Microsoft Windows ME™ OS includes a feature called "System Restore". That system is essentially a snapshot system, and only backs up files related to the OS and installed applications (not user files).

[0009] A current practice of maintaining computers is to image the hard drive of a computer while in a working state. If the computer becomes unstable, or if undesirable content appears on the computer, the computer's drive is restored using the earlier made image. This practice is lacking in that all changes made following the image creation are wiped off the system when the

[0010] Also, some applications are not provided with an uninstall program. To de-install those applications an administrator is required to know where the application files and settings reside in the system, and remove them manually.

[0011] It is therefore apparent that much time and money is expended in the administration of applications on computing platforms, and thus there is a need for a way to ease the installation and de-installation of applications, and prevent application conflicts.

BRIEF SUMMARY OF THE INVENTIONS

[0012] The inventions relate generally to computer systems having facilities for providing virtual portions of file systems and configuration settings to applications. More particularly, the inventions relate to computer systems that provide a layer organization for files and configuration settings that can be overlaid on top of an operating system. Detailed information on various example embodiments of the inventions are provided in the Detailed Description below, and the inventions are defined by the appended claims.

BRIEF DESCRIPTION OF THE DRAWINGS

[0013] Figure 1 illustrates components of a layering computer system at a conceptual level.

[0014] Figure 2 illustrates an operation of a layering computer system at a conceptual level.

[0015] Figure 3 illustrates components of a particular layering computer system.

[0016] Figure 4 illustrates components of a layering computer system at simple organizational level.

[0017] Figure 5 shows a simplified method for performing read file system operations using a layered computing system.

[0018] Figure 6 shows a simplified method for performing write file system operations using a layered computing system.

[0019] Figure 7 shows a simplified method for performing registry operations using a layered computing system.

[0020] Reference will now be made in detail to some embodiments of the inventions, example of which are illustrated in the accompanying drawings.

DETAILED DESCRIPTION

[0021] General Concepts

[0022] For the purpose of simplifying the discussion herein, an example computing device may be referenced. That device is a conventional personal computer or workstation having a CPU, memory, display, keyboard, mouse, and at least one fixed disk. It will be apparent to one of ordinary skill in the art that the concepts disclosed herein may apply equally to other computing systems that are not personal computers, for example diskless workstations, headless workstations or servers, and embedded systems. Herein it is contemplated that the inventions may be applied to these and other computing systems, both existing and yet to be, using the methods and principles disclosed herein.

[0023] Likewise the discussion below speaks of Registries and registry settings, which are specific to Microsoft Windows™ operating systems. It will be recognized that registry settings are merely configuration for the operating system and applications installed to a computing device, accessible through a system-wide API. The meaning of registries and registry settings is therefore extended to future Windows operating systems and operating systems other than Windows, where equivalent structures and access facilities exist thereon.

[0024] In the discussion below, the words "enabled" and "activated" are used interchangeably to describe layers that are active or enabled on a layering computing system. Likewise, the words "disabled" and "deactivated" may be used to describe layers that are not enabled or active.

[0025] Provided in one aspect of the invention are application layers which are isolated from other applications on a computer. In that aspect, an application layer may be defined to be a group of files in combination with any associated application configuration stored to operating system files. An application of a layered system may be an application in the most commonly used meaning, such as word processors, browsers, system tools, games, and the like, or may extend to other software installed to a host providing an environment, such as a graphical user

environment or shell. It will be seen that isolating application files and configuration in a layer provides several benefits, including the ability to delete, disable, and enable applications in a simple way and to provide a barrier between applications which may use conflicting configuration or library files. The use of a layering system may therefore enhance the stability, reliability, usability and security of a computing system.

[0026] A layered system introduces a new concept of organizing data from disparate sources and presenting a virtual view of that data to an operating system and a user. This permits the real data to be much more logically organized while still presenting to the operating system and the user an expected view and access of that data. In a sense, a layer is a higher order storage unit. Because a layer can be managed as a unit for the purposes of exporting, importing, enabling, disabling, and so on, a computer system and user data can be managed with a greater degree of flexibility and reliability, also with improved security. As changes to a layered system are made, the changes are organized while being written, rather than tracking the changes made. By doing this both a speed penalty and the dedication of large amounts of storage for images and changes are avoided.

[0027] Depicted in figure 1 are components of a layering computer system at a conceptual level. A base operating system 110 forms a platform with which applications can be run and files can be accessed in file systems. Base operating system 110 further has registry settings, globally available to applications for reading and writing. The system has libraries 108 for executing the functions of the operating system including operating file systems and registries, and other operating system functions. Tied into libraries 108 are layering system libraries and/or software 106 which intercept file system and registry accesses from applications 104. As accesses are received from applications 104, the layering system software 106 performs computations to determine whether the accesses should be permitted to continue to the base operating system 110, or should be redirected to layer information 112, the information relating to and the contents of files and registry settings. A layer manager application 100 may be provided to permit control and configuration of the layering system software 106 through a management API and library 102.

[0028] Depicted in figure 2 is the operation of a layering computer system at a conceptual level. An application 200 is running on a layered computing system. This computing system contains a base file system 206, and two layers labeled "A" and "B", 204 and 202 respectively. In this example layer B has priority over layer A, which in turn has priority over the base file system. A first file access 208 is made by application 200. The layered computing system determines the owner of the file being accessed. Finding an entry for file access 208 in layer B, the corresponding file in layer B is opened and returned to the application. The file access of

208 might also correspond to files in layers A or the base file system, however layer B is determined to be the owner as it has priority over layer A and the base. Another file access 210 is made by application 200. The computing system does not, however, find a corresponding entry in layer B. An entry is found in layer A, which has priority over the base file system. Again, if a file existed in the base file system corresponding to the file access, it would be accessed because layer A is found to be the owner with priority. The computing system is not able to find corresponding entries in layers A or B for file access 212, so that access is made to the base file system.

[0029] In figure 4 components of a layering computer system at simple organizational level are shown. A computing device includes a processor 400, which may also have peripheral devices attached such as memory, input devices or output devices as desired. Processor 400 interacts with one or more storage devices 402, providing storage for the processor. On storage 402 is a base operating system 408 and applications 410. A number of layers 404a-n are also contained on storage 402, each having applications 406a-n.

[0030] In larger aspects, a layer may be defined to be a set of file system and registry changes, that combination forming an organizational unit that may be managed by layered system software. In some cases, a layer need not contain registry changes, but only changes to one or more file systems. In those cases it may be desirable to limit support in the layered system software to files of file systems. A layer definition may include layer properties and settings, layer inclusive files, references to those files, registry settings and locations, and a manifest or directory those file and registry references.

[0031] References may be made inherent, if desired, by locating files and registry settings in a structure that mirrors a real underlying file system. Such a mirroring system may be organized in a common directory, with one subdirectory per defined layer, each containing a mirrored directory structure of the underlying file system.

[0032] An exported layer will contain all of the layer-included information bundled in a transportable archive. Exported layers may be further bundled into groups, which is especially useful for layers that rely on other layers, such as layers of a hierarchy or peer layers. For systems that utilize a mirror structure of an underlying file system, it may be desirable to hide the mirror structure from applications, except perhaps a manager application, so as to prevent accidental data modification, loss, or meddling.

[0033] A layer intending to isolate an application has stored thereon the files and directory

structure of the application's installation. When that layer becomes mounted (or enabled), those application files and directories are shadowed or overlaid over the regular operating system file system. Shared libraries (such as DLLs), system accessible configuration (such as registry entries), and version control are managed by the layering subsystem, optionally using an internal database. Though each layer is a separate and individual entity within the host OS, the application files, data, and system accessible configuration are presented as if they resided in their respective ordinary locations. Thus an application stored in a layer appears to the host OS as if it were installed in the ordinary fashion with the expected functionality.

[0034] For example, suppose a layer existed in a Windows OS environment that specified that in C:\windows there should be a file called winfile.exe. Suppose that this file did not reside in the real C:\windows directory. When the layer is not active, a file listing of C:\windows does not show a winfile.exe. When the layer becomes active, the layering system merges (or overlays) the real listing of C:\windows and the file list described in the layer. In this example, applications (and thereby a user) would see all of the files in the real C:\windows directory and winfile.exe. Registry values in a layer may be handled in a similar manner.

[0035] Shown in figure 5 is a simple method for performing read file system operations using a layered computing system. A loop is entered beginning at step 500. Execution halts in step 502 pending the receipt of a read request. A determination is then made in step 504 as to whether or not the file reference of the request is maintained in an enabled layer. To perform that determination all the layers on the system are generally examined for a virtual file corresponding to the file reference of the request. If no enabled layer contains such a virtual file, step 506 executes in which the usual read operation is executed using the file reference of the request. Otherwise, an owner layer is identified in step 508. For example, if two enabled layers contain a virtual reference to a file, one will take priority over the other and be identified as the owner layer. Step 510 then executes, in which a virtual file reference is determined that corresponds to the file reference of the read request. That virtual file reference might be an offset and length for a storage device in some systems, a pathname at a mirrored location in other systems, or other reference. Afterward, the read operation is executed using that virtual file reference in step 512.

[0036] Figure 6 shows a simple method for performing write file system operations using a layered computing system. A loop is entered beginning at step 600. Execution halts in step 602 pending the receipt of a write request. A determination is then made in step 604 as to whether or not the file reference of the request should be captured to an enabled layer. That determination may be made, for example, by noting the state of the system software is in a capture state, and in some circumstances by noting the PID of the calling application and

parents. If no enabled layer is configured for capture, step 606 executes in which the usual write operation is executed using the file reference of the request. Otherwise, a capture layer is identified in step 608. Step 610 then executes, in which a virtual file reference is determined that corresponds to the file reference of the write request. That virtual file reference might be an offset and length for an unused portion of a storage device in some systems, a pathname at a mirrored location in other systems, or other reference. Afterward, the write operation is executed using that virtual file reference in step 612.

[0037] The read and write operations spoken of in the discussion of figures 5 and 6 may be performed on some systems through an open() call. A read request, for example, might be a call to open() with a pathname as a file reference and "r" as an option. Likewise, a write request might be a call to open with "w" or "+" as an option. In either case, a file handle is returned which would correspond either to a true file reference (if the file reference is not managed in a layer) or to a virtual file reference (if the file reference is managed in at least one layer). That file handle will continue to be used in data read and write operations, and thus the data will be delivered to and from the correct system locations. Other systems may use other equivalent methods of opening, reading and writing, and applicable using the methods described herein.

[0038] Figure 7 shows a simple method for performing registry operations using a layered computing system. The method begins at step 700, following which a pause is executed at step 702 until a request for registry setting operation is received. When a registry setting request is received, step 704 executes in which a determination is made as to whether or not the request is to be captured to an enabled layer. If not, step 706 is executed in which a usual registry function is called, as if layering were not present in the system. Otherwise, step 708 is performed, in which a destination layer is identified. Step 710 tests the request for a registry entry creation request. If a creation request was received, step 712 executes in which a virtual registry entry is created in the destination layer. Otherwise step 714 is performed, testing for a registry entry deletion request. If positive, step 716 is executed in which either a virtual registry entry is deleted, if the entry exists in a single layer, or a delete entry is made in the virtual registry of the destination layer signifying that the registry entry should not appear while that layer is enabled. If the request is neither a create or delete request, step 718 is performed testing for a set registry entry request. If positive, step 720 executes creating a virtual setting in the destination layer.

[0039] As in the above example, layers may contain file and registry deletion references. Those references may be used where a layer specifies the absence of a file or registry setting, whereby a specified file or registry setting will appear to be absent from the computing system

only when the layer is enabled.

[0040] The use of a layering system provides several advantages. If applications are stored individually in layers, interactions between application files may no longer occur due to conflicting shared libraries (DLLs), as each application 'sees' only its own installed libraries first, followed by libraries in the base operating system, those base libraries optionally preceded by libraries from other layers if desired. Applications captured in a layer may be safely and completely uninstalled by simply removing the layer from the host computing system. Different versions of an application may be stored as layers on a single computer; the user may select a desired version by enabling the particular layer. A layering system may also extend the file systems of the OS beyond physical limits if layers are stored on separate disk partitions or remote file systems. If layering is used for a group of installed applications, the computing system may be restored to a "virgin" or checkpoint state by removing one or a group of application layers. The transfer of applications between similar computing systems can be simplified, in that the transfer may be done simply by moving the layer containing the application. The bundling of an application and user files into a layer provides a package that may be compressed or encrypted and transported conveniently. Using a layering system application vendors can provide 'pre-installed' applications as layers on CD-ROM or other media, those applications being pre-tested and guaranteed to work with a high reliability. A layer also provides a convenient container to limit access to an application, for example for time limited use or license key access.

[0041] In preferred systems, the enablement and disablement of layers is performed through a system call. The system drivers control the access of applications to the file system through the enabled layers, generally without requiring significant access to the system disk or other storage. In those systems the installation and de-installation of an application can be as simple as enabling or disabling a containing layer, without requiring the installation or removal of the applications files from a hard disk. In those systems, time consuming snapshot utilities become unnecessary.

[0042] In a preferred system, layering only applies to files located to fixed disks and network drives, each layer spanning one or more fixed disks. In those systems removable disks should not generally be layered, as a layer generally pertains to the persistent files and configuration required to operate an application or user environment. It is expected that under most circumstances user files should be permitted to be saved to a floppy disk or CD-RW, for example, so a user can transport his files to another computer. Likewise, areas on a fixed disk may also be reserved for user or other files to be excluded from layering, for example a "my documents" directory, as desired.

[0043] In some systems it will be advantageous to distinguish layers into a "read-only" and "read-writable" portions, the read-only portion containing files and configuration as originally installed and the read-writable portion containing additions, deletions and modifications to the original installation. In some circumstances these layers may be referred to as the install portion (read-only) and the user (read-write) section. A read-writable portion may be global to all users of a computer. Alternatively a read-writable portion may be provided for each user of a computer, each read-writable portion being protected from access by other users.

[0044] Some systems provide a multi-user environment providing a facility for an administrator to designate layers accessible to individual users and another facility to automatically enable layers on user login and disable layers after a user has logged off. In those systems an administrator may provide layers accessible to all users or some users. Other layers may be provided accessible only to an individual user. In a subset of those systems a writable layer is provided for each user, providing data protection and isolation between users.

[0045] A single layer having a read-only and a read-writable portion is equivalent to two layers, one of which is write protected. In alternate systems, read-only and read-writable layer portions are individual peer layers; those layer definitions containing a reference to the accompanying peer layer.

[0046] In layered systems layers may be stacked on top of each other, with the real file system at the bottom of the stack. If files of the same name and location exist in multiple layers, or in the base file system, rules can be provided whereby the layered system can determine which file to present to an application. In some systems, layers include dependency information. That dependency information may include a list of layer identifiers which are required to be enabled when a particular layer is enabled. Dependencies may be asserted when a layer is created, by recording the layers enabled on a layered system at the time of layer creation. The layering system software may automatically enable all dependent layers when a particular layer is enabled.

[0047] For ease of configuring and managing a layering system, a manager application may be provided. The manager application permits an administrator or user to control the presentation of applications and data on a system, as well as other functions. A manager application may have facilities for importing and exporting layers, using a standard layer archive format. That archive format will advantageously be compressed, and may use standard

archiving formats, for example those used by 'zip' or 'tar' type applications. A manager application provides a logical place to contain a facility for changing layered system software settings. A manager application might provide a viewer to view information about a layer. Likewise, a layer editor may be provided to edit certain layer information as desired. An editor might also be provided whereby registry settings and files can be added, removed, or changed in a layer. A facility for selecting, enabling, and disabling layers and layer groups may also be provided. Likewise, a facility for defining and editing layer groups may be included, as well as layer dependency information. A facility for deleting and installing layers may also be provided in that manager application. That application may also include an interface to cause layered system software to enter and exit capture modes.

[0048] It may also be desirable to provide a startup layer enablement function, whereby the computing system starts up a group of layers based on layer configuration. This will be especially helpful where it is desired not to provide users with non-layered access to the underlying file system and registry, for example in public settings.

[0049] It may optionally be desired to include variable handling with regard to file system paths and registry paths. The location of a file or registry setting specified in a layer may include one or more variables, so as to permit relocation of that object. A variable may be denoted in many ways, for example by surrounding the variable name with percent "%" characters. The source of some variable names and values may be from the environment. For example, Windows operating systems set the "WINDIR" environment variable to the location of the OS system subtree, for example C:\windows. Including the WINDIR variable in a path may permit files of a layer to be moved from one Windows system to another, especially if the OS system subtree resides in different locations on the computers. Other variable values may be supplied at runtime, for example a "CURRENTUSER" variable. In that example, the CURRENTUSER variable is set to a user's login name while that user is logged in. One use of the CURRENTUSER variable is to provide a layered file reference for a global file that appears in each user's profile directory. Yet other variable names and values may be stored in a layer definition. A manager application may provide editing facilities for changing those layer-defined variables, and for editing the pathnames of virtual files.

[0050] Layer Creation Modes

[0051] Layer creation modes may be provided in a layered system to create new layers through a "capture" operation. A capture operation is generally started and ended, and uses the layering software to intercept operations that install, delete, rename or modify files and configuration such as a registry. If the layering system supports layers having both a readable

and read-writable portion, the capture operation may record changes to the readable portion; that readable portion becoming effectively locked when the capture operation is ended. During the capture operation changes made by the installation procedure do not affect the base system, but are rather recorded to the new layer.

[0052] A first layer creation mode is simply called "capture" mode. When that mode is enabled, all operations by any application to create, modify or delete files are entered into a layer. This mode is especially helpful in situations where it is desirable to create a new layer for one or more applications to be installed to the computing system. In an example of a capture mode operation on a Windows platform, a user first enables capture mode. The user then executes an application installation program. During the install, all of the applications shared DLLs, registry entries, and .ini files that would be directed to the Windows system directories become trapped in the capture layer. Application files that would be placed on file systems managed by the OS are also redirected into the layer. All of the captured data is held separate from the regular OS either locally or remotely in a data file, hard disk partition, or some other container.

[0053] A second layer creation mode is referred to as "capture by PID" mode. That mode is similar to "capture" mode, with the difference being that only changes made by a particular process ID (PID) or one of its child PIDs are captured.

[0054] A third layer creation mode is called "delete capture" mode. This mode may be thought of as the inverse of "capture" mode. Delete capture mode is intended to track all of the file system and registry deletions that occur and place those files and registry entries into a new layer. The software (driver) is hooked into the system so that operations that delete, rename, or modify file system or registry so they can be copied to the capture layer before they are modified. This mode may be particularly helpful to create a layer of an already installed application. The user enters "delete capture" mode, following which the user activates the application's deinstallation program. As the application's uninstall program removes files and registry settings, they are copied to the new layer. When the uninstall is complete, the user exits delete capture mode. At that time the application does not exist in the regular file system or registry, but can be activated by the user as it appeared before the uninstall operation by activating the newly created layer.

[0055] A fourth layer creation mode is called "delete capture PID" mode. That mode operates in similar fashion to delete capture mode, with the difference that only changes made by a particular PID and child PIDs are tracked, rather than system-wide changes.

[0056] A system supporting layering need not implement a capture mode if an alternate layer delivery mechanism is provided, for example a layer import operation or a simple file or file system copy.

[0057] Use: Application Installation Generator

[0058] Many application installer programs have the ability to create an application install via a "capture" or "snapshot" process. This process typically involves comparing the state of the computer system before and after an application install and generating the install information based on the differences. In a system supporting layers, an application may be captured as outlined above, creating an installation layer. Because changes are tracked as they occur, no state comparison needs to be done, saving time. In addition, it is usually recommended that the "capture" operation be performed on a "clean" or "virgin" system, so the capture process can capture all the necessary system changes (i.e. won't miss changes due to application pieces being left over from prior installations.) This requires the user to reinstall the operating system to get the system into the desired clean state. A layered system may be made clean by disabling all layers created during installation capture procedures (assuming all install operations have occurred under capture operations.) After capture of an installation layer, that layer can be used to install the application at another computer supporting layers, or the information can be extracted from the layer to provide a file manifest for other installation programs.

[0059] Use: Software Installation/Uninstallation

[0060] Layers can be advantageously used to provide an installation for an application that is relatively easy to uninstall. A software vendor builds an application CD (or other media), first using a capture mode to record a layer of the application installation. That layer is then exported to a file, which file is then combined with an installation program for the layering system software, for example to a compact disc. The compact disc will contain an installation program, which for example might be called 'setup'. The setup program operates first to install the layering system software, and then import the layer exported to the compact disc into the destination system. At that point, the subject application is then installed to the destination system, but isolated in a layer. Because it is isolated, it is protected from corruption from other applications or meddling, and thus it remains in a known and reliable state, potentially reducing the number of technical support calls.

[0061] It is probably desirable to include a banner screen advertising the layering system software product and providing contact information for product inquiry and support. It may also

be desirable to include a layer manager application with the layering system software to allow a user to enable and disable the application layers or other installed layers, but that is not necessary for a simple demonstration product.

[0062] As the application is used, it may be desired to record changes to the virtual file system into the writable portion of a layer. Alternatively, it may be desirable to record some user files to the underlying file system so those files could be retained if the application layer was deinstalled or removed, as might be the case for word processing files, CAD files, etc. The software installer may be given the option to record the software installation of an application layer into a readable-only portion, so the user cannot inadvertently or otherwise damage the application installation.

[0063] At some point, it may be desired to remove the application. To do so, the user removes the layer from his computer, which deinstalls the application and any files or changes made to the virtual file system. Uninstalling the layering system software is optional, as the presence of that software would not adversely affect the use of the destination system.

[0064] Through that method, software creators may create a demo version of their software. These versions might be used to give the end user experience with a software product before a purchase is made. The advantage of removing changes to the virtual file system is significant, as many applications do not uninstall cleanly and leave residual files and changes.

[0065] Optionally, functionality might be built into the layering system software that disables the application layer after a period of time. After such a disabling, a user would have the option of removing the application layer, or purchasing a license for use of the application. The license would presumably be accompanied with a license key or other authentication feature verifiable by the layering system software.

[0066] In another alternative configuration, an application layer is never transferred to a resident fixed disk, but rather remains resident on the vendor product, compact disc or otherwise. In that configuration the application layer can only be activated if the vendor product is readable in a media drive, and little or no space is taken on resident file systems for the application installation.

[0067] Use: Secure Applications

[0068] Applications can be protected from unauthorized access through the use of a layered system. In a first situation, it is desired to protect application files from viewing and copying, for

which one solution is described. The layering system software of a suitable system has an additional feature by which a layer may include an authentication key or token. Any application of the computing system desiring to open files within the layer must supply a token to the layering system software before access is allowed. The PID of an authenticating application may be tracked so that only one authentication step is required. The application layer may additionally be encrypted, the layering system software performing decryption and encryption steps at runtime as the application layer is accessed. That system is advantageous in that only the data of a particular application need be encrypted, reducing the complexities of bootstrapping into an encrypted file system and modifying system applications to support encrypted system files.

[0069] In that system authenticating applications will have access to the application files, but not applications not having a valid authentication token. The authenticating applications can be constructed such that limited access is permitted to the application files, as desired by the programmer. For example, an application may store a license key to one of the application's files. If access were permitted to that file, an unscrupulous user could copy that license key to a second computer providing illicit access to the application software stored thereon. The authenticating layered system software is installed to the computer, and an application layer is constructed and installed to the computer, that layer encrypted using a key constructed with information specific to the computer, for example a volume label or Ethernet MAC address. A second application installed to the computer, for example Windows Explorer, cannot view the application layer files because it does not possess the correct authentication key. A user is therefore prevented from copying or otherwise accessing the application files, providing security for the software vendor.

[0070] In a second situation, it is desirable to protect the software from execution by unauthorized individuals. In that system, the layering system software has a facility for authenticating a user before enabling a layer.

[0071] Use: Secure Base OS

[0072] In some circumstances it is desirable to regularly restore a computer to a 'virgin' state. This is sometimes done, for example, by Internet cafes and college computer laboratories, or other systems typically used by untrusted users. The computers are regularly reverted back to a known good state to ensure that users have a stable and working system free viruses and from potential interference and security risks. The computer restoration is often performed by writing an image containing all files of the computer to the computer's hard drive made earlier. A layered computing system can serve better.

[0073] To use a layered computing system in a first system protection and restoration mode, an administrator first installs a base OS and applications to a computer. The administrator then causes the computer to enter a capture mode, by which all further changes are recorded to a layer and not to the underlying file systems and OS resources. Users may then use the computer in an unrestricted fashion, including installing applications, using the Internet, and even passing (inadvertently) viruses to the computers. At the end of the day (or other period), the administrator ends the capture mode, deletes the layer, and re-enters capture mode. All changes made by users are then wiped and the computer is restored to its base state.

[0074] A second system protection and restoration mode provides that user data may be retained between sessions. To use this mode, an administrator first installs a base OS and applications to a computer. The administrator then sets up the computer such that when a user logs in (or otherwise starts a user session), a user layer is enabled. While in a user session, changes are recorded to that user's layer. When the user logs out, the user layer is disabled and retained for future use. That layer may be used repeatedly for multiple user sessions, the state of the user's files and configuration being maintained between sessions. Applications specific to a user might be installed to a user layer providing an easy way to control what applications are available to a particular user.

[0075] If desired, a user layer may be stored to a network server and retrieved and stored as required to any computer of a computer farm, so the user may access his files and state anywhere in the computer group. That layer may provide files and directories overlaid on the existing fixed disk file systems, as opposed to locations referencing new network mount points. On a Windows directory structure files located to a layer might appear under C:, as opposed to being presented as a new and separate volume (e.g., D:).

[0076] Alternatively, a layer might be configured to store files presented as a virtual volume. For example, a user might store his data in a virtual volume accessible under P:, and backup his user data by backing up the user layer. That layer could also be protected using encryption and authentication by including appropriate facilities in the layered system software. In a variation on that system, a layer might be configured to present files in a subdirectory on an existing volume which could be at any level in a file system hierarchy.

[0077] A layering system might also be used to provide a "bottomless" virtual storage device. In that system an auxiliary storage device is presented by layering system software as being available on a main local storage device (e.g. C:). The auxiliary storage device might be an additional local fixed disk, network drive or file system, or other storage resource. The auxiliary

storage device space is effectively added to the existing space, providing a way to expand an existing fixed disk that has become filled. The auxiliary storage device may optionally be hidden to the system. This system is advantageous in that no repartitioning is required, no backup and restore option is required, and no uninstall/reinstall operation of applications is required.

[0078] Use: Heirarchical Layers

[0079] Multiple layers can advantageously be used. In one layer heirarchy, one layer represents a user type and a writable layer contains the user's changes. For example, a company has 500 computers. These computers all have the same base software installed, which may be only the OS. The company then has layers defined for different types of users, for example secretaries, engineers and accountants. The secretarial layer contains word processing, spreadsheet, and other secretarial applications. The engineering layer contains software development tools, CAD tools, and other engineering related applications. The accounting layer includes accounting software. In addition, each user may have a personal layer, which contains an individual's changes on top of the type layer and base system.

[0080] If a user causes his computer to fail, an administrator can restore the computer by disabling the user's personal layer. If a computer is to be transferred from engineering to accounting, the administrator removes the engineering type layer and installs the accounting type layer. Using the above exemplified heirarchical layer organization can simplify the administration of a large number of workstations in a company or other organization.

[0081] Other Uses

[0082] Another use for a layering system is to have layers that represent different environments on a system. For example, a user could have an Office and a Gaming layer, each providing an environment with it's own icons and menus.

[0083] In another use, multiple versions of a software product are installed on a computer, each isolated in a layer. A user may enable a particular layer and use the corresponding version of the software product without having to de-install and re-install the application. This use may be especially helpful where an older version of a software product supports a function desired but not supported in a newer version, for example, the importation of older word processing file formats. This use would also be useful to software product testers, who in the course of testing verify software functionality against multiple development versions. In that use the repeated unstalling and reinstalling or cleaning and reinstalling operations are avoided.

[0084] Example Implementation

[0085] Example systems are provided of an application layering system under a 32-bit Microsoft Windows architecture, such as Windows 95, 98, NT, 2000, and XP. In those system a layering system is formed by adding several files to the stock Windows operating system, those files including a runtime library FSLLIB32.DLL, a compression/archiving library, and an FSLX driver which is either an FSLX.VXD driver (for 95/98/ME based platforms) or an FSLX.SYS driver (for NT based platforms). The addition of those files is performed using an installation program. The example layering system provides a user with the ability to contain third party application installations into a "file system layer" or FSL. The example system provides the modes of "capture", "capture by PID", "delete capture", and "delete capture PID".

[0086] Depicted in figure 3 are components of the example layering computer system. An operating system 314 is installed to a computing device, that operating system having subsystems for handling a registry 316 and a file system 318. An FSL system driver is installed "on top" of the operating system 314 in order to have first processing priority for registry and file system accesses. An FSL management application 300 provides an administrator an interface to interact with the FSL system driver 312, change its configuration, and make changes to layers. An FSL API library 306 provides a convenient interface for the management application 300 to attach to the FSL system driver 312. At certain times, FSL management application 300 provides notices to the Windows Explorer 302 notifying that application that the contents of a mounted file system have changed. Other applications 304 may interact with the system, performing read and write operations to the file system and registry, through the FSL system driver 304. A compression library 310 may be provided to compress layer information, especially for layer archives exported by the system.

[0087] A "lazy thread" is utilized to perform low priority tasks. That thread wakes up occasionally to delete layers that are marked for deletion and write out delete lists that have changed. The execution of the lazy thread may be deferred for a short time if the system is busy.

[0088] In the example systems there is a key in the registry under HKEY_LOCAL_MACHINE\SYSTEM called FSLogic\FSL where registry settings describe each layer and its settings. The SYSTEM portion of the registry is used because it is available very early in the boot cycle. Each layer has the properties outlined in the following table:

Property/Value	Meaning/Function
Active	Non-zero indicates that the layer is enabled
ActiveOnStart	Non-zero indicates the layer should be enabled when the FSLX driver loads.
FileRedirect	The path to the location in the file system that contains the file system virtual files.
MajorVersion	The major version of the layer format.
MinorVersion	The minor version of the layer format.
Peer	The name of the peer layer.
ReadOnly	Non-zero indicates that the layer is read only, or the readable portion of a peer layer combination.
RegRedirect	Path to the location that contains the virtual registry settings for the layer.
Type	Layer type.
ShouldDelete	Non-zero value indicates that the layer should be deleted. This value is read by the lazy thread to know if the layer should be deleted.

[0089] Also under HKEY_LOCAL_MACHINE\SYSTEM under a key called fsldr is kept all registry information contained in each layer. Under fsldr there is further a key for each layer defined in the system. Under each layer key each of the HCC, HCR, HCU, HLM, and HU keys are present. These keys correspond to HKEY_CURRENT_CONFIG, HKEY_CLASSES_ROOT, HKEY_CURRENT_USER, HKEY_LOCAL_MACHINE, and HKEY_USERS respectively. The structure of the registry under these keys mimics the regular structure of the system registry.

[0090] When a layer is active, all of the keys and values for the layer are overlaid on the normal registry. For example, a layer "TEST" is defined on a system and has a registry entry "HKEY_LOCAL_MACHINE\SYSTEM\fsldr\TEST\HLM\Software\XYZCorp". When that layer becomes active, the following key would appear in the registry:
"HKEY_LOCAL_MACHINE\Software\XYZCorp".

[0091] The FSLX.SYS and its counterpart FSLX.VXD operate to intercept key file system and registry calls and manipulate the results to create the appearance that virtual files and registry settings contained in the layer definitions exist in the real file system and real registry. When requests come that access virtual files or virtual registry settings, these requests are redirected by the FSLX driver to the proper locations in the layer. The FSLX driver also accepts IOCTLs from FSLLIB32.DLL that control the state of the driver. The following table outlines a set of IOCTL commands available through the FSLX driver:

IOCTL	Description
Version Query	Returns the driver version.

<i>IOCTL</i>	<i>Description</i>
Begin Capture	Causes the driver to enter "Capture" mode.
End Capture	Causes the driver to exit "Capture" mode.
Begin Delete Capture	Causes the driver to enter "Delete Capture" mode.
End Delete Capture	Causes the driver to exit "Delete Capture" mode.
Activate Layer	Activates a specified layer.
Deactivate Layer	Deactivates a specified layer.
Rename Layer	Notifies the driver that a layer has been renamed.

[0092] For each read or write request to a file system or registry, an owner layer is determined. The owner layer is determined by a sequence of steps. First, if the driver is in Capture mode, the owner layer is the layer being captured. Second, if the driver is in PID Capture mode, and if the PID of the requesting process is the PID being captured or a child PID of the PID being captured, the owner layer is the layer being captured. Lastly, if the driver is not in capture mode, and if the PID of the requesting process is a PID of an executable whose executable file is in a layer, the owner layer is the layer where the executable file resides.

[0093] Because multiple layers can be active at the same time and those layers may contain entries that overlap, rules are defined to determine the order layers are considered by the driver. Different modes require different search rules. If the system is in capture mode, the owner layer is defined to be the layer specified as the capture layer. Otherwise, the owner layer is defined to be the layer that a particular process started from, as may be determined by traversing upward the PID parent/child chain. For example, suppose layer A contained B.EXE. When B.EXE executes, it results in process C running. The owner layer for process C is then layer A.

[0094] When the FSLX driver loads, the following is performed: (1) all mutexes and lists are initialized, (2) a device is created used for API DLL communications, (3) a symbolic link that allows for the device object's access from Win32 programs is made, (4) all of the file system entry points are hooked in, (5) the drives to be redirected (C:, D:, etc.) are hooked in, (6) all of the Registry entry points are hooked in, (7) the lazy thread is started.

[0095] The FSLX driver uses the following structures and hooks the following entry points in the file system and Registry system code:

[0096] Structures used:

[0097] FSLX_DELETE_ENTRY_REMOVE: Holds information about an entry on a delete list that may be removed later, for which all necessary information will not be available at the time of removal.

[0098] FSLXDELETIONCANDIDATE: Holds information about a file that should be later marked as deleted.

[0099] PFSLXOPENREGHANDLE: Holds information about all currently open registry handles.

[0100] FSLX_PFO_ENTRY: Holds information about an open directory, the information including a pointer to the file object, a handle to the directory, and the directory path.

[0101] FSLX_RENAME_ENTRY: Holds information about a rename operation that is used to create a delete entry.

[0102] FSLXREGOPENKEY: Holds information about an open key in a layer, including a handle to the key.

[0103] SH_RET_ENTRY: Holds the name of a file. These file names may have already been returned in a query routine. This structure is retained to ensure the same name is not returned more than once if the same file exists in multiple redirection areas.

[0104] FSLXSHADOWHANDLE: Holds information about an open directory handle. Among other things, it may contain a list of FSLX_PFO_ENTRYs that correspond to directories in applicable layers.

[0105] FSLXSHADOWKEY: Holds information about an open registry key. Among other things, it may contain a list of FSLXREGOPENKEY structures that correspond to keys in applicable layers.

[0106] File System Calls:

[0107] IRP_MJ_CLEANUP: If there is an FSLX_DELETE_ENTRY_REMOVE structure associated with the parameter referenced File Object, free it. If there is an FSLXDELETIONCANDIDATE structure associated with the parameter referenced file object, add a delete entry for the file and free the structure.

[0108] IRP_MJ_CLOSE: Free the FSLXSHADOWHANDLE structure associated with the parameter referenced File Object by: (1) removing the shadowHandle from the list, (2) getting a pointer to the shadowhandle using the File Object, (3) decrement the reference count of the shadowHandle, (4) if the reference count is greater than zero, return success, otherwise (5) free the originalPath member of the shadowHandle, (6) for each FSLX_PFO_ENTRY: remove the entry from the list, free the file path, dereference the File Object, close the directory handle, and free the FSLX_PFO_ENTRY structure, (7) for each SH_RET_ENTRY: remove the entry from the list and free the name and structure, (8) free the search string, and (9) free the structure.

[0109] IRP_MJ_CREATE: Get the full file name and full parent directory path for the request. Determine if the File Object represents a file or a directory. If the File Object represents a directory, determine if it represents the root directory. Check to see if this is a reentrant call for which the SL_OPEN_TARGET_DIRECTORY bit in currentIrpStack->Flags should be set. If this is a reentrant create, get the shadowHandle object for this File Object, increment the reference count on the shadowHandle if there is one, and return. Determine the owner layer. If the path of the file being opened is in a redirected area, and if the file that is being created is on the delete list, create and fill in an FSLX_DELETE_ENTRY_REMOVE structure and return. The completion routine for that operation checks to see if the create was successful and, if so, removes the delete entry from the delete list. Check to see if the create is for a *.Config or a *.Manifest file. If it is, set a flag, for which at the completion of this routine if the return code is STATUS_OBJECT_PATH_NOT_FOUND the return code is changed to STATUS_OBJECT_NAME_NOT_FOUND. If the request is for a directory, do (1) if a shadowHandle already exists for the parameter referenced File Object, increment it's reference count, (2) if a shadowHandle does not exist, create one with all entries initialized to default values, and for each layer that contains a corresponding directory or delete entries that correspond to the directory, create an FSLX_PFO_ENTRY entry. Determine if the parameter referenced request should be redirected: (1) if the request is a write request and capture mode is enabled, do (a) make sure the parent directory is in the layer being captured, (b) if the parameter referenced request is to a file and if a delete entry exists for the file, create an FSLX_DELETE_ENTRY_REMOVE structure so that the delete entry can be removed if this

operation is successful, (c) if the parameter referenced request is to a file and if a delete entry does not exist for the file, use the standard search order to locate and copy any existing file to the writable portion of the layer being captured, and (d) redirect the create to the writable portion of the layer being captured and return; (2) if no layers have the directory and it is an open (not a create), don't redirect and return from the function call; (3) if there is no owner layer, do: (a) if the request is a write request, don't redirect and return from the function call, (b) if the request is a read request, find a first file by iterating through each layer in the search path, and redirect to that file unless the file is on a delete list; (4) if an owner layer can be identified, and if the request is a write request: (a) make sure the directory path exists in the writable section of the owner layer, (b) if the parameter referenced request is to a file, and if a delete entry exists for the file, create an FSLX_DELETE_ENTRY_REMOVE structure so that the delete entry can be removed upon function call completion, (c) if the parameter referenced request is to a file, and if no delete entry exists for the file, use the standard search order to locate and copy any existing file to the writable portion of the layer being captured, and (d) redirect the writable portion of the layer being captured and return; and (5) if an owner layer can be identified, and if the request is a read request, find a first file by iterating through each layer in the search path, and redirect to that file unless the file is on a delete list. If the file that is being opened is on the delete list, return STATUS_OBJECT_NAME_NOT_FOUND. If the open is being performed with the FILE_DELETE_ON_CLOSE flag, and if the parameter referenced file is a file that should be protected from delete, (1) clear the FILE_DELETE_ON_CLOSE flag, and (2) create an FSLXDELETIONCANDIDATE structure, later used in the completion routine to add a delete entry for the file. Return a value that indicates success or failure.

[0110] IRP_MJ_CREATE: (completion routine) If the create operation is being canceled, free the shadowHandle if one exists, free any existing FSLXDELETIONCANDIDATE and return. If the create operation failed, free any existing shadowHandle and FSLXDELETIONCANDIDATE and return. If an FSLX_DELETE_ENTRY_REMOVE exists, use it to remove the delete entry from the delete list.

[0111] IRP_MJ_DIRECTORY_CONTROL: If the minor function code is IRP_MN_QUERY_DIRECTORY, (1) get the shadowHandle for the File Object, (2) if there is no shadowHandle, return, (3) if the root directory is being enumerated, do not return "." or ".." entries, (4) enumerate the corresponding directories in each layer and the real directory. Use SH_RET_ENTRY structures to make sure duplicate entries are not returned.

[0112] IRP_MJ_SET_INFORMATION: If the FileInformationClass is

FileDispositionInformation, if the file is being deleted, and if it is a file that should be protected from deletion, create an FSLXDELETIONCANDIDATE structure to be used in the completion routine to add a delete entry for the referenced file. Otherwise, if FileInformationClass is FileRenameInformation, do the following: (1) if the requested operation is a rename operation on a protected file that should succeed, copy the source file to the writable section of the owner layer and create a delete list entry for the source file, (2) if the requested operation is a rename operation on an unprotected file, perform the rename operation and create an FSLX_RENAME_ENTRY entry for the source file.

[0113] IRP_MJ_SET_INFORMATION: (completion routine) If FileInformationClass is FileRenameInformation, and if there is an FSLX_RENAME_ENTRY, use the contained information to create a delete entry for the source file of the rename operation. If FileInformationClass is FileDispositionInformation, do: (1) if the operation was successful and the file was deleted, get the FSLXDELETIONCANDIDATE structure, and if the deleted file was not in the writable section of the owner layer, cancel the deletion, (2) if the operation was successful and the delete operation was canceled, remove any existing FSLXDELETIONCANDIDATE, or (3) if the operation was unsuccessful, and if a deletion was being attempted, remove any existing FSLXDELETIONCANDIDATE.

[0114] Registry Calls:

[0115] RegCloseKey: If this call is re-entrant, pass the call parameters to the OS. Since all NtClose calls come through this hook and not just RegCloseKey calls, make sure that this call is a close for a registry handle. If not, pass the call parameters to the OS. Get the shadowKey structure. If there exists a shadowKey, (1) free the shadowKey and all FSLXREGOPENKEY structures by closing the handle to the key and freeing the structure, and (2) if the main key handle has not been closed, close it. If there is no shadowKey, close the handle. Remove any PFSLXOPENREGHANDLE.

[0116] RegCreateKey: If this call is re-entrant, pass the call parameters to the OS. If requesting in a redirected part of the registry, pass the call parameters to the OS. Get the PID of the caller. If there is a delete entry corresponding to the requested create operation, (1) create a new key in the writable section of the owner layer, (2) if unable to create the key, return an error, (3) change the disposition to REG_CREATED_NEW_KEY, (4) create a new shadowKey structure for the created key, (5) determine the owner layer for the key, (6) if there

is an owner layer (a) allocate a new FSLXSHADOWKEY structure and initialize with default values and (b) create an FSLXREGOPENKEY entries for applicable layers, (7) if the key does not exist in the base registry, but does in one or more layers, create a user mode handle to be returned to the calling application, and (8) remove the delete entry. Otherwise if there is no delete entry corresponding to the requested create operation, continue. Create a shadowKey structure. Determine the owner layer for the key. If there is an owner layer (1) allocate a new FSLXSHADOWKEY structure and initialize with default values, and (2) create FSLXREGOPENKEY entries for applicable layers. If the key does not exist in the base registry but it does in one or more layers, create a user mode handle to be returned to the calling application. If the key can be opened (not created), set the disposition to REG_OPENED_EXISTING_KEY, create a new PFSLXOPENREGHANDLE and return. If creation of a key in the writable section of an owner layer is successful, do: (1) set the disposition to REG_CREATED_NEW_KEY, (2) create a PFSLXOPENREGHANDLE, and (3) return. If the error code from the creation attempt was STATUS_OBJECT_PATH_NOT_FOUND, return STATUS_OBJECT_PATH_NOT_FOUND. If a key was not created in the writable section of an owner layer, attempt to create the key in the base registry, create a PFSLXOPENREGHANDLE, and return.

[0117] RegDeleteKey: If this call is re-entrant, pass the call parameters to the OS. Otherwise, if there is an owner layer, do: (1) if the key has child keys, return STATUS_ACCESS_DENIED, or (2) if the key has no child keys, create a delete entry for the key. If there is no owner layer, do: (1) if there is a shadowKey, delete the key from the base registry and add delete entries to all layers, or (2) if there is no shadowKey, delete the key from the base registry.

[0118] RegDeleteValueKey: If this call is re-entrant, pass the call parameters to the OS. Otherwise, if there is an owner layer, create a delete entry for the value. If there is no owner layer, delete the value from the real registry and create delete entries for all applicable layers.

[0119] RegEnumerateKey: If this call is re-entrant, pass the call parameters to the OS. Otherwise, if there is a shadowKey, (1) enumerate through the read registry and applicable layers, (2) store state information in the shadowKey. Do not return duplicate entries. If there is no shadowKey, pass the call parameters to the OS.

[0120] RegEnumerateValueKey: If this call is re-entrant, pass the call parameters to the OS.

Otherwise, if there is a shadowKey, (1) enumerate through the read registry and applicable layers, (2) store state information in the shadowKey. Do not return duplicate entries. If there is no shadowKey, pass the call parameters to the OS.

[0121] RegFlushKey: If this call is re-entrant, pass the call parameters to the OS. If there is a shadowKey, flush the real registry key and all applicable layer keys. Otherwise, pass the call parameters to the OS.

[0122] RegOpenKey: If this call is re-entrant, or if the key is in the redirection area of the registry, pass the call parameters to the OS. Otherwise, get the caller's PID. If there is a delete entry for this open, return STATUS_OBJECT_NAME_NOT_FOUND. Create a shadowKey. Try to identify an owner layer. If an owner layer can be identified, (1) allocate a new FSLXSHADOWKEY structure initialized with default values, (2) create FSLXREGOPENKEY entries for applicable layers, and if a key does not exist in the base registry but it does in one or more layers, create a user mode handle to be returned to the calling application. If the open operation was successful, create a PFSLXOPENREGHANDLE.

[0123] RegQueryKey: If this call is re-entrant, pass the call parameters to the OS. If there is no shadowKey and the request is of class "KeyNameInformation", get the key name and if it is the name of a redirect key, change it to the base name. If there is a shadowKey and there is a delete entry found for this key, return STATUS_OBJECT_NAME_NOT_FOUND. If there is a shadowKey and there is not a delete entry for this key, query the real registry key and all applicable layer keys. Depending on the class of query, combine the results and return them to the user.

[0124] RegQueryValueKey: If this call is re-entrant, or if there is no shadowKey, pass the call parameters to the OS. If there is a delete entry for this value, return STATUS_OBJECT_NAME_NOT_FOUND. Otherwise, if there is a shadow key, use the standard search order to find the value to return.

[0125] RegSetValueKey: If this call is re-entrant, or if there is no owner layer, pass the call parameters to the OS. Otherwise, set the value in the writable portion of the owner layer. If the setting operation was successful, remove any delete entry for the value and return.

[0126] In the example systems the FSLLIB32.DLL runtime library provides an API that may be used by other applications to manage the layered system and communicate with the FSLX driver, and further provides system management function implementations. That library includes functions to load and unload the FSLX driver, identify version information for itself and FSLX driver; begin and end Capture mode; begin and end Delete Capture mode; import and export layers; create, delete, rename and merge layers; activate and deactivate layers; get layer information; enumerate layers; enumerate the files of a layer; enumerate the registry entries of a layer; manipulate the registry entries of a layer; enable and disable layers; set and unset an "active on start" layer property, create and delete layer groups; enumerate layer groups; add and remove layers from layer groups; verify system integrity; enumerate layer variables; create and delete layer variables; and delete the writable portion of a layer and create a new, empty writable portion. A discussion of the individual exported functions follows with greater specificity, using C language prototypes:

<i>Function</i>	<i>Description</i>
FSLActivate(PTCHAR fsIName)	Validates the fsIName against defined layers. If corresponding layer or group is defined, get information. If fsIName corresponds to a group, recursively call FSLActivate for each layer in the group. Communicates with FSLX driver via an IOCTL to activate the layer. Notifies the Windows Explorer that classes may have changed. For each virtual directory contained in the newly activated layer, notify the Windows Explorer that the directory contents have changed. Applications in the layer that are specified to be run on system startup (in win.ini, registry, startup folder, etc.) are started. Return a value indicating success or failure.

<i>Function</i>	<i>Description</i>
FSLAddLayerToGroup(PTCHAR fsIName, PTCHAR groupName)	Verifies that both the specified layer and group are defined. Creates a subkey under the group key with the name of the layer, adding the layer to the group. Return a value indicating success or failure.

<i>Function</i>	<i>Description</i>
FSLAddVariable(PTCHAR fsIName, PTCHAR varName, PTCHAR varValue)	Verifies the specified layer is defined. Open the variables key for the specified layer. Set a registry value using the provided varName and varValue. Return a value indicating success or failure.

Function	Description
FSLCreate(PTCHAR fsIName, BOOL createPeer)	Verifies the specified layer is not defined. Create a layer definition with default values. Create the layer redirection area in the file system(s). If createPeer is true, recursively call FSLCreate for the peer with createPeer set to FALSE, and set the peer entries in the layers to point to each other. Return a value indicating success or failure.

Function	Description
FSLCreateGroup(PTCHAR groupName)	Validates groupName. If the group already exists, return an error. Create a new group named groupName under the group key HKEY_LOCAL_MACHINE\SYSTEM\FSLogic\groups) Return a value indicating success or failure.

Function	Description
FSLDeactivate(PTCHAR fsIName, BOOL force, PDWORD pPid)	Validate fsIName, and get information about the corresponding layer or group. If fsIName corresponds to a group, recursively call FSLDeactivate for each layer of the group. If fsIName corresponds to a layer, communicate with the FSLX driver through an IOCTL to deactivate the layer. <u>If the FSLX driver returns an error that there is a PID running from this layer and force is true, kill the PID corresponding to pPid.</u> Return a value indicating success or failure.

Function	Description
FSLDelete(PTCHAR fsIName, BOOL deletePeer, BOOL force, PDWORD pPid)	Validates fsIName. If the corresponding layer does not exist, or if the corresponding layer has not been deactivated, return an error. If deletePeer is TRUE, recursively call FSLDelete with the name of the peer layer, with deletePeer set to FALSE. Mark the layer as deleted. Remove the fsldr registry branch for the corresponding layer. Remove the layer from any group entries. Return a value indicating success or failure.

Function	Description
FSLDeleteGroup(PTCHAR groupName)	Validates groupName. Deletes the group key and any subkeys or values. Return a value indicating success or failure.

Function	Description
FSLDeletePeer(PTCHAR fsIName, BOOL force, PDWORD pPid)	Validates fsIName. Finds the peer for fsIName. Calls FSLDelete using the found peer name. Return a value indicating success or failure.

Function	Description
FSLDeleteVariable(PTCHAR fslName, PTCHAR varName)	Validates fslName. Delete any variable/value pair from the layer's variables key. Return a value indicating success or failure.

Function	Description
FSLEnable(PTCHAR fslName, BOOL bEnable)	Validate fslName, and get information about the corresponding layer or group. If fslName corresponds to a group, recursively call FSLEnable using the same bEnable for each layer of the group. If fslName corresponds to a layer, set the enabled value of the corresponding layer based on bEnable. Return a value indicating success or failure.

Function	Description
FSLEndCapture(PTCHAR fslName)	Validate fslName. Communicates with FSLX driver through an IOCTL call to cause the driver to exit capture mode. Notifies Windows Explorer that classes may have changed. For each directory contained in the newly activated layer, Windows Explorer is notified that the directory contents have changed. Return a value indicating success or failure.

Function	Description
FSLEXP (PTCHAR fslName, PTCHAR archivePath, BOOL replacelfExists, PTCHAR errorStr, void (__stdcall *RTInfoFunc)(PFSL_IMP_EXP pImpexp), BOOL bInitialCall)	Validate fslName, and get information about the corresponding layer or group. If bInitialCall is TRUE, perform a number of initialization steps including (1) validating the archivePath, (2) testing for the existence of an archive file in the archivePath directory, (3) if the replacelfExists flag is FALSE, returning an error if an archive file already exists in the archivePath directory, (4) if the replacelfExists flag is TRUE, deleting an archive file located in the archivePath directory, (5) if fslName corresponds to a layer having a peer layer, recursively calling FSLEXP once for both the corresponding layer and the peer layer with bInitialCall set to FALSE, followed by closing the archive file. Otherwise, if fslName corresponds to a layer group, perform a number of steps including (1) for each layer of the group, recursively calling FSLEXP for each layer and any existing peer layer to each layer with bInitialCall set to FALSE, (2) storing the group name in the archive, (3) placing a version number in the archive, and (4) closing the archive file. If bInitialCall is FALSE and fslName corresponds to a layer, perform the steps of (1) creating a new archive file if it has not yet been created, (2) opening the archive file, (3) exporting the fslrdr portion of the registry of the layer to a new file, (4) exporting the layer definition in the system registry to a new file, (5) creating a file designating the name of the layer, (6) adding all of the created files in the previous three steps plus the files in the redirection area of the file systems of the layer to the archive, (7) placing a version number in the archive, (8) closing the archive file, and (9) removing the exported registry files and layer name designation file. Return a value indicating success or failure.

Function	Description
FSLFindClose (HANDLE hFindFile)	Call FindClose (of the WIN32 API) using hFindFile. Return a value indicating success or failure.

Function	Description
FSLFindCloseGroup (PFSL_FIND *groupFind)	Close the registry key in groupFind. Return a value indicating success or failure.

Function	Description
FSLFindCloseLayer (PFSL_FIND *fslFind)	Close the registry key in fslFind. Return a value indicating success or failure.

Function	Description
FSLFindCloseLayerInGroup (PFSL_FIND *fslFind)	Close the registry key in fslFind. Return a value indicating success or failure.

Function	Description
FSLFindCloseVariable(PFSL_FIND *find)	Close the registry key in find. Return a value indicating success or failure.

Function	Description
FSLFindFirstFile(LPCTSTR fslName, LPCTSTR lpFileName, LPWIN32_FIND_DATA lpFindFileData)	Validate fslName. Generate a search string including the redirection area of the layer and lpFileName. Call FindFirstFile (WIN32 API) on the redirect search string. Return a value indicating success or failure.

Function	Description
FSLFindFirstGroup(PFSL_FIND *groupFind, PTCHAR groupName)	Open the parent key in the registry where all group names are stored (HKEY_LOCAL_MACHINE\SYSTEM\FSLLogic\groups). Set the index in groupFind to 0. Find the first group name. Return a value indicating success or failure.

Function	Description
FSLFindFirstLayer(PFSL_FIND *fslFind, PTCHAR fslName, BOOL includePeers)	Open HKEY_LOCAL_MACHINE\SYSTEM\FSLLogic\fsl. Store the handle to the key in the fslFind structure. Set the index in the fslFind structure to 0. Set includePeers in the fslFind structure to the value of includePeers. Get the first layer name from the registry (layer names are subkeys of HKEY_LOCAL_MACHINE\SYSTEM\FSLLogic\FSL). <u>If a layer is marked for deletion, go to the next layer.</u> Skip peer layers if includePeers is FALSE. Return a value indicating success or failure.

Function	Description
FSLFindFirstLayerInGroup(PFSL_FIND *fslFind, PTCHAR groupName, PTCHAR fslName)	Open the group registry key under HKEY_LOCAL_MACHINE\SYSTEM\FSLLogic\groups. Set the index in fslFind to 0. Get the first layer name from the registry. Return a value indicating success or failure.

Function	Description
FSLFindFirstVariable(PFSL_FIND *find, PTCHAR varName)	Open the variables registry key under the layer definition key. Set the index in find to 0. Find the first value (<u>is this a var name or var value?</u>). Return a value indicating success or failure.

Function	Description
FSLFindNextFile(HANDLE hFindFile, LPWIN32_FIND_DATA lpFindFileData)	Call FindNextFile (WIN32 API). Return a value indicating success or failure.

Function	Description
FSLFindNextGroup(PFSL_FIND *groupFind, PTCHAR groupName)	Increment the index in groupFind. Read the next group name from the registry. Return a value indicating success or failure.

Function	Description
FSLFindNextLayer(PFSL_FIND *fslFind, PTCHAR fslName)	Increment the index in the fslFind structure. Read the next layer name from the registry. Skip layers marked for deletion. If the includePeers field in fslFind is FALSE, skip peer layers. Return a value indicating success or failure.

Function	Description
FSLFindNextLayerInGroup(PFSL_FIND *fslFind, PTCHAR fslName)	Increment the index in fslFind. Read the next layer name from the group key. Return a value indicating success or failure.

Function	Description
FSLFindNextVariable(PFSL_FIND *find, PTCHAR varName)	Increment the index in find. Find the next value (<u>is this a var name or var value?</u>). Return a value indicating success or failure.

Function	Description
FSLGetDriverVersion(PDWORD pdMajVersion, PDWORD pdMinVersionstruct)	Communicates to the FSL Driver via an IOCTL call to determine the FSL driver's major and minor version numbers. Sets pdMajVersion and pdMinVersion to the major and minor version numbers of the FSL driver. Return a value indicating success or failure.

Function	Description
FSLGetInfo(PTCHAR fsIName, PFSL_INFO *pInfo)	Validate the fsIName. Set structure pointed to by pInfo to zero. Copy the layer name into the structure. If fsIName corresponds to a group, (1) set blsGroup in pInfo to TRUE, and (2) look at all the layers in the group and set enabled, active, and activeOnStart flags of the pInfo structure appropriately. Read the active, enabled, activeOnStart, majorVersion, minorVersion, type, and peerName values from the registry and set the corresponding flags of the pInfo structure. Return a value indicating success or failure.

Function	Description
FSLGetVersion(PDWORD pdMajVersion, PDWORD pdMinVersionstruct)	Sets pdMajVersion and pdMinVersion to the major and minor version numbers of the FSLX driver. Return a value indicating success or failure.

Function	Description
FSLGetVariable(PTCHAR fsIName, PTCHAR varName, PTCHAR varValue)	Read the value named by varName from the specified layer's variables key into varValue. Return a value indicating success or failure.

Function	Description
FSLImport(PTCHAR archivePath, BOOL replaceIfExists, PTCHAR errorStr, void (__stdcall *RTInfoFunc)(PFSL_IMP_EXP pImpexp))	Verify the archivepath (the archivepath being the full pathname to the file). Open the archive file. Check the version numbers against what is supported by the FSLX driver (i.e. driver version number > archive version number), returning an error if unsupported. Extract the files that contain the layer and group names. Create each group. For each layer to be imported, perform the following: (1) if a layer of the same name already exists and if replaceIfExists is FALSE return an error, otherwise delete the existing layer, (2) extract all pertinent information for the layer from the archive, (3) delete the file that indicates the layer name, (4) import the registry fsldr branch for the layer, (5) import the layer definition, (5) mark the layer as enabled, and (6) delete the layer registry information files. Close the archive. Return a value indicating success or failure.

Function	Description
FSLInitSystem(void)	Verify the FSL system: (1) make sure HKEY_LOCAL_MACHINE\SYSTEM\FSLLogic\FSL exists, (2) make sure major and minor version registry value are created, (3) make sure default file system redirection path and registry redirection path registry values are set, (4) make sure HKEY_LOCAL_MACHINE\SYSTEM\fsldr exists, and (5) make sure C:\fsldr exists. Read default file system redirection path. Read default registry redirection path. Return a value indicating success or failure.

Function	Description
FSLIsGroup(PTCHAR name)	Validate the name. Determine if name is a valid group by attempting to open the group key under HKEY_LOCAL_MACHINE\SYSTEM\FSLLogic\groups. Return a value indicating success or failure.

Function	Description
FSLLoadDriver(void)	Verify the FSL system: (1) make sure HKEY_LOCAL_MACHINE\SYSTEM\FSLLogic\FSL exists, (2) make sure major and minor version registry value are created, (3) make sure default file system redirection path and registry redirection path registry values are set, (4) make sure HKEY_LOCAL_MACHINE\SYSTEM\fsldr exists, and (5) make sure C:\fsldr exists. Loads the driver if it is not loaded. Notifies Windows Explorer via SHChangeNotify that the C:\fsldr directory has changed. Return a value indicating success or failure.

Function	Description
FSLRegCloseKey(HKEY hKey)	Close the registry key. Return a value indicating success or failure.

Function	Description
FSLRegCopyKey(HKEY srcKey, PTCHAR srcKeyName, HKEY destParentKey, BOOL overwrite, BOOL removeAfterCopy)	Create a new key name under the destination parent key. If the key already existed under the destination parent and overwrite is FALSE, and if copying the values and subkeys from the source would overwrite any values or subkeys in the destination return FALSE. Otherwise, copy the subkeys and values to the destination. If removeAfterCopy is TRUE, delete the registry source key with all of its subkeys and values. Return a value indicating success or failure.

Function	Description
FSLRegCopyValue(HKEY srcKey, LPCTSTR lpValueName, HKEY destKey, BOOL overwrite, BOOL removeAfterCopy)	If the value already exists under destKey and overwrite is false, return an error. Read the source value and write that value to the destination. If removeAfterCopy is TRUE, remove the source value (<u>what about the source key?</u>) Return a value indicating success or failure.

Function	Description
FSLRegCreateKeyEx(HKEY hKey, LPCTSTR lpSubKey, DWORD Reserved, LPTSTR lpClass, DWORD dwOptions, REGSAM samDesired, LPSECURITY_ATTRIBUTES lpSecurity Attributes, PHKEY phkResult, LPDWORD lpdwDisposition)	Create a registry path to the layer's redirection area using the layer's redirect path, its name, and lpSubKey. Create the key in the redirection area. Return a value indicating success or failure.

Function	Description
FSLRegDeleteKey(HKEY hKey, LPCTSTR lpSubKey)	Remove the key and all subkeys and values. Return a value indicating success or failure.

Function	Description
FSLRegDeleteValue(HKEY hKey, LPCTSTR lpValueName)	Delete the specified value. Return a value indicating success or failure.

<i>Function</i>	<i>Description</i>
FSLRegEnumKeyEx(HKEY hKey, DWORD dwIndex, LPTSTR lpName, LPDWORD lpcbName, LPDWORD lpReserved, LPTSTR lpClass, LPDWORD lpcbClass, PFILETIME lpftLastWriteTime)	Enumerate the specified key. Return a value indicating success or failure.

<i>Function</i>	<i>Description</i>
FSLRegEnumValue(HKEY hKey, DWORD dwIndex, LPTSTR lpValueName, LPDWORD lpcbValueName, LPDWORD lpReserved, LPDWORD lpType, LPBYTE lpData, LPDWORD lpcbData)	Enumerate the specified value. Return a value indicating success or failure.

<i>Function</i>	<i>Description</i>
FSLRegOpenKeyEx(PTCHAR fsName, HKEY hKey, LPCTSTR lpSubKey, DWORD ulOptions, REGSAM samDesired, PHKEY phkResult)	Create a registry path to the layer's redirect area using the layer's redirect path, its name, and lpSubKey. Open the key in the redirection area. Return a value indicating success or failure.

Function	Description
FSLRegQueryValueEx(HKEY hKey, LPTSTR lpValueName, LPDWORD lpReserved, LPDWORD lpType, LPBYTE lpData, LPDWORD lpcbData)	Query the value specified. Return a value indicating success or failure.

Function	Description
FSLRegSetValueEx(HKEY hKey, LPCTSTR lpValueName, DWORD Reserved, DWORD dwType, CONST BYTE *lpData, DWORD cbData)	Set the specified value. Return a value indicating success or failure.

Function	Description
FSLRemoveLayerFromGroup(PTCHAR fslName, PTCHAR group)	Verify that the group exists, and that the layer is a member of the group. Remove the layer from the group by deleting the key with the layer's name from the group key. Return a value indicating success or failure.

Function	Description
FSLResetPeer(PTCHAR fslName, BOOL force, PDWORD pPid)	Get the peer name for this layer (writable section of the layer). Get information about the peer. make sure the peer is deactivated. Delete the peer. Create the peer. Point the layer and the new peer layer at each other by setting their peer values in the registry. If the named layer is active, activate the new peer layer. Return a value indicating success or failure.

Function	Description
FSLSetActiveOnStart(PTCHAR name, BOOL bActiveOnStart)	Verify the name corresponds to an existing layer or group. Get information about the named layer or group. If the name corresponds to a group, recursively call FSLSetActiveOnStart for each layer in the group. Otherwise, set the activeOnStart value for the layer to bActiveOnStart. Return a value indicating success or failure.

<i>Function</i>	<i>Description</i>
FSLSetLayerInfo(PTCHAR name, PTCHAR fileRedirect, PTCHAR regRedirect, DWORD *pType, DWORD *pReadOnly, PTCHAR peerName)	Verify that the name corresponds to a layer. Open the registry key that contains the layer definition. If fileRedirect is specified, set the value of the proper registry value. If regRedirect is specified do: (1) set the value of the proper registry value, (2) create the specified redirect path, (3) create the redirect root keys (HLM, HCU, HU, HCC, and HCR). If type is specified, set the value of the proper registry value. If readOnly is specified, set the value of the proper registry value. If peerName is specified, set the value of the proper registry value. Return a value indicating success or failure.

<i>Function</i>	<i>Description</i>
FSLStartCapture(PTCHAR fslName, BOOL bTrack, DWORD dPid)	Validates fslName to make sure it is a valid layer name (legal characters, etc.) Communicates to the FSL Driver via an IOCTL to put it into Capture mode. Notifies Windows Explorer that classes may have changed. For each directory contained in the newly activated layer, Windows Explorer is notified that the directory contents have changes. Applications in the layer that are specified to be run on system startup are started (there are several places where these can be specified: win.ini, registry, startup folder, etc.) Return a value indicating success or failure.

<i>Function</i>	<i>Description</i>
FSLUnloadDriver(BOOL force)	All active layers are deactivated. Unloads the FSLX driver. Notifies Windows Explorer via SHChangeNotify that the C:\fslrdr directory has changed. Return a value indicating success or failure.

<i>Function</i>	<i>Description</i>
FSLVerifyBaseSystem(void)	Make sure HKEY_LOCAL_MACHINE\SYSTEM\FSLLogic exists. Put the current major and minor version into majorVersion and minorVersion values. Put the default File System redirection path in a DefaultFileRedirect value. Put the default Registry redirection path in a DefaultRegistryRedirect value. Make sure HKEY_LOCAL_MACHINE\SYSTEM\fslrdr exists. Make sure fslrdr exists at the root of all file systems that will be redirected. Return a value indicating success or failure.

[0127] Each of the above functions returns a value of the type "FSLLIB32_API DWORD _stdcall" indicating success or failure. In the above functions, the TCHAR variable type changes depending on the compilation options. If compiled for Unicode, a TCHAR is a 16 bit entity, otherwise it is an 8 byte char. A BOOL may be represented by a single bit, but is often defined to be a word so as to permit efficient word alignment according to the processor

architecture. A DWORD is normally a 32-bit integer. And an LPCTSTR is a long pointer to a constant array of TCHARs.

[0128] In the example systems, on each file system volume (C:, D:, etc.) included in the system there is an fsldr directory at the root of the volume. This directory contains file system information for each of the defined layers. Under the fsldr directory directories that correspond to each layer are maintained. Under each of those layer directories is a directory that represents the drive letter. Under each of those letter directories the contained directories and file structures mimic the regular structure of the containing drive. When a layer is active all of the directories and files defined for the layer are overlaid on the normal file system. For example, the directory "C:\fsldr\TEST\c\XYZCorp" is defined under a "TEST" layer. When the "TEST" layer is active, the directory "c:\XYZCorp" appears on the C: drive to all applications running under that layer, and optionally under other layers depending on the implementation details.

[0129] While the present systems and methods have been described and illustrated in conjunction with a number of specific configurations, those skilled in the art will appreciate that variations and modifications may be made without departing from the principles herein illustrated, described, and claimed. The present invention, as defined by the appended claims, may be embodied in other specific forms without departing from its spirit or essential characteristics. The configurations described herein are to be considered in all respects as only illustrative, and not restrictive. All changes which come within the meaning and range of equivalency of the claims are to be embraced within their scope.

CLAIMS

What is claimed is:

1. A set of computer readable media containing computer instructions for operating a layered computing environment, the set of computer readable media comprising at least one medium upon which is stored the computer instructions executable by a computing system to achieve the functions of:
 - (i) receiving from applications a read request for a read operation to a file system, the read request containing a file reference appropriate to the file system organization;
 - (ii) a first determining whether or not the file reference is maintained in at least one enabled layer;
 - (iii) if in the first determining a file reference is found not to be maintained in at least one enabled layer, causing the read operation to execute using the file reference of the read request;
 - (iv) if in the first determining a file reference is found to be maintained in at least one enabled layer, identifying an owner layer from the set of enabled layers;
 - (v) following the identifying an owner layer, identifying a virtual read reference utilizing information contained in the layer; and
 - (vi) following the identifying a virtual read reference, causing the read operation to execute using the virtual read reference.
2. A set of computer readable media according to claim 1, wherein the computer instructions are further executable to achieve the function of decrypting encrypted information referenced by the virtual read reference.
3. A set of computer readable media according to claim 1, wherein the computer instructions are further executable to achieve the functions of:
 - (vii) receiving from applications a write request for a write operation to a file system, the write request containing a file reference appropriate to the file system organization;
 - (viii) a second determining whether or not the file reference is a reference to a write operation to be captured in an enabled layer;
 - (ix) if in the second determining a file reference is determined not to be a reference to a write operation to be captured to an enabled layer, causing the write operation to execute using the file reference of the write request;
 - (x) if in the second determining a file reference is determined to be a reference to a write operation to be captured to an enabled layer, identifying a capture layer;
 - (xi) following the identifying a capture layer, creating a virtual write reference

corresponding to the file reference of the write request; and

(xii) following the creating a virtual write reference, causing the write operation to execute using the virtual write reference.

4. A set of computer readable media according to claim 3, wherein the computer instructions are further executable to achieve the function of copying an existing file to the capture layer prior to causing the write operation to execute if the write operation is a modify operation.

5. A set of computer readable media according to claim 3, wherein the computer instructions are further executable to achieve the function of encrypting the information to be written prior to the execution of the write operation.

6. A set of computer readable media according to claim 3, wherein the computer instructions are further executable to achieve the functions of:

(xiii) receiving from applications requests to create, delete, and set the value of a registry setting;

(xiv) following receipt of a request to create, delete or set the value of a registry setting, a determining whether or not the registry operation of the request is to be captured to an enabled layer;

(xv) acting on a request to create a registry setting, and on a determination that a registry setting is not to be captured to an enabled layer, causing the registry setting operation to execute in the base system registry location;

(xvi) acting on a request to create a registry setting, and on a determination that a registry setting is to be captured to an enabled layer, identifying a registry creation destination layer;

(xvii) following said identifying a registry creation destination layer, causing the registry setting to be created virtually in the registry creation destination layer;

(xviii) acting on a request to delete a registry setting, and on a determination that a registry setting is not to be captured to an enabled layer, causing the registry setting deletion operation to execute in the base system registry location;

(xix) acting on a request to delete a registry setting, and on a determination that a registry setting is to be captured to an enabled layer, identifying a registry deletion destination layer;

(xx) following said identifying a registry deletion destination layer, causing the registry setting to be deleted virtually in the registry deletion destination layer;

(xxi) acting on a request to set a registry setting, and on a determination that a registry setting is not to be captured to an enabled layer, causing the registry setting operation to execute in the base system registry location;

(xxii) acting on a request to set a registry setting, and on a determination that a registry setting is to be captured to an enabled layer, identifying a registry setting destination layer; and

(xxiii) following said identifying a registry setting destination layer, causing the registry setting to be created virtually in the registry setting destination layer.

7. A set of computer readable media according to claim 6, wherein the computer instructions are further executable to achieve the function of:

(xiii) receiving management commands through an applications programmer interface.

8. A set of computer readable media according to claim 7, wherein the computer instructions are further executable to achieve the functions of:

(xiv) receiving a management command through an applications programmer interface to import a specified layer archive; and

(xv) import the specified layer archive.

9. A set of computer readable media according to claim 7, wherein the computer instructions are further executable to achieve the functions of:

(xiv) receiving a management command through an applications programmer interface to enable a layer group; and

(xv) enable each layer of the specified layer group.

10. A set of computer readable media according to claim 7, wherein the computer instructions are further executable to achieve the functions of:

(xiv) receiving a management command through an applications programmer interface to disable a layer group; and

(xv) disable each layer of the specified layer group.

11. A set of computer readable media according to claim 7, wherein the computer instructions include a layer manager application.

12. A set of computer readable media according to claim 6, wherein the computer instructions are further executable to achieve the functions of:

(xiii) examining the installed layers for a configuration element, that element specifying for each layer whether or not the layer is to be enabled on system initialization; and

(xiv) enabling those layers having configuration elements specifying layer enablement on system initialization.

13. A computing system supporting a layered computing environment, comprising:

a processor;

a storage device group accessible by said processor, said storage device group containing at least one storage device operable to contain operating system files, applications and layers;

one or more layers stored to said storage device group;

instructions stored to said storage device group, said instructions being further executable by said processor to achieve the functions of:

- (i) receiving from applications a read request for a read operation to a file system, the read request containing a file reference appropriate to the file system organization,
- (ii) a first determining whether or not the file reference is maintained in at least one enabled layer,
- (iii) if in the first determining a file reference is found not to be maintained in at least one enabled layer, causing the read operation to execute using the file reference of the read request,
- (iv) if in the first determining a file reference is found to be maintained in at least one enabled layer, identifying an owner layer from the set of enabled layers,
- (v) following the identifying an owner layer, identifying a virtual read reference utilizing information contained in the layer,
- (vi) following the identifying a virtual read reference, causing the read operation to execute using the virtual read reference,
- (vii) receiving from applications a write request for a write operation to a file system, the write request containing a file reference appropriate to the file system organization,
- (viii) a second determining whether or not the file reference is a reference to a write operation to be captured in an enabled layer,
- (ix) if in the second determining a file reference is determined not to be a reference to a write operation to be captured to an enabled layer, causing the write operation to execute using the file reference of the write request,
- (x) if in the second determining a file reference is determined to be a reference to a write operation to be captured to an enabled layer, identifying a capture layer,
- (xi) following the identifying a capture layer, creating a virtual write reference corresponding to the file reference of the write request, and
- (xii) following the creating a virtual write reference, causing the write operation to execute using the virtual write reference.

14. A computing system according to claim 13, wherein the computer instructions are further executable to achieve the functions of:

- (xiii) receiving from applications requests to create, delete, and set the value of a registry setting;

(xiv) following receipt of a request to create, delete or set the value of a registry setting, a determining whether or not the registry operation of the request is to be captured to an enabled layer;

(xv) acting on a request to create a registry setting, and on a determination that a registry setting is not to be captured to an enabled layer, causing the registry setting operation to execute in the base system registry location;

(xvi) acting on a request to create a registry setting, and on a determination that a registry setting is to be captured to an enabled layer, identifying a registry creation destination layer;

(xvii) following said identifying a registry creation destination layer, causing the registry setting to be created virtually in the registry creation destination layer;

(xviii) acting on a request to delete a registry setting, and on a determination that a registry setting is not to be captured to an enabled layer, causing the registry setting deletion operation to execute in the base system registry location;

(xiv) acting on a request to delete a registry setting, and on a determination that a registry setting is to be captured to an enabled layer, identifying a registry deletion destination layer;

(xx) following said identifying a registry deletion destination layer, causing the registry setting to be deleted virtually in the registry deletion destination layer;

(xxi) acting on a request to set a registry setting, and on a determination that a registry setting is not to be captured to an enabled layer, causing the registry setting operation to execute in the base system registry location;

(xxii) acting on a request to set a registry setting, and on a determination that a registry setting is to be captured to an enabled layer, identifying a registry setting destination layer; and

(xxiii) following said identifying a registry setting destination layer, causing the registry setting to be created virtually in the registry setting destination layer.

15. A computing system according to claim 13, wherein the computer instructions are further executable to achieve the function of maintaining a delete list.

16. A computing system according to claim 13, wherein the computer instructions are further executable to achieve the function of executing a lazy thread.

17. A computing system according to claim 13, wherein the computer instructions are further executable to achieve the function of decrypting encrypted information referenced by a virtual read reference.

18. A computing system according to claim 13, wherein:

the computer instructions are further executable to achieve the function of (xiii) receiving management commands through an applications programmer interface; and
the computer instructions include a layer manager application.

19. A set of computer readable media according to claim 18, wherein the computer instructions are further executable to achieve the functions of:

(xiv) receiving a management command through an applications programmer interface to enable or disable a layer group; and

(xv) enable or disable each layer of the specified layer group.

20. A method of handling requests to read and write files in a layered computing system, the method comprising the steps of:

(i) receiving from applications a read request for a read operation to file systems, the read request each containing a file reference appropriate to the file system organization;

(ii) for received read requests, performing a first determining whether or not the contained file references are maintained in at least one enabled layer;

(iii) if in the first determining file references are found not to be maintained in at least one enabled layer, causing the corresponding read operations to execute using the file references of the read requests;

(iv) if in the first determining file references are found to be maintained in at least one enabled layer, for each file reference identifying an owner layer from the set of enabled layers;

(v) following identifying an owner layer, identifying a virtual read reference utilizing information contained in the identified layer for each file reference;

(vi) following the identifying a virtual read reference, causing read operations to execute using the virtual read references;

(vii) receiving from applications write requests for write operations to file systems, the write requests each containing a file reference appropriate to the file system organization;

(viii) for received write requests, performing a second determining whether or not the contained file references are references to write operations to be captured in an enabled layer;

(ix) if in the second determining file references are determined not to be references to write operations to be captured to an enabled layer, causing the write operations to execute using the file references of the write requests;

(x) if in the second determining file references are determined to be references to write operations to be captured to an enabled layer, identifying a capture layer for each file reference;

(xi) following the identifying a capture layer, creating virtual write references corresponding to the file references of the write requests; and

(xii) following the creating virtual write references, causing the write operations to execute using the virtual write references.

PAGE INTENTIONALLY LEFT BLANK

1/5

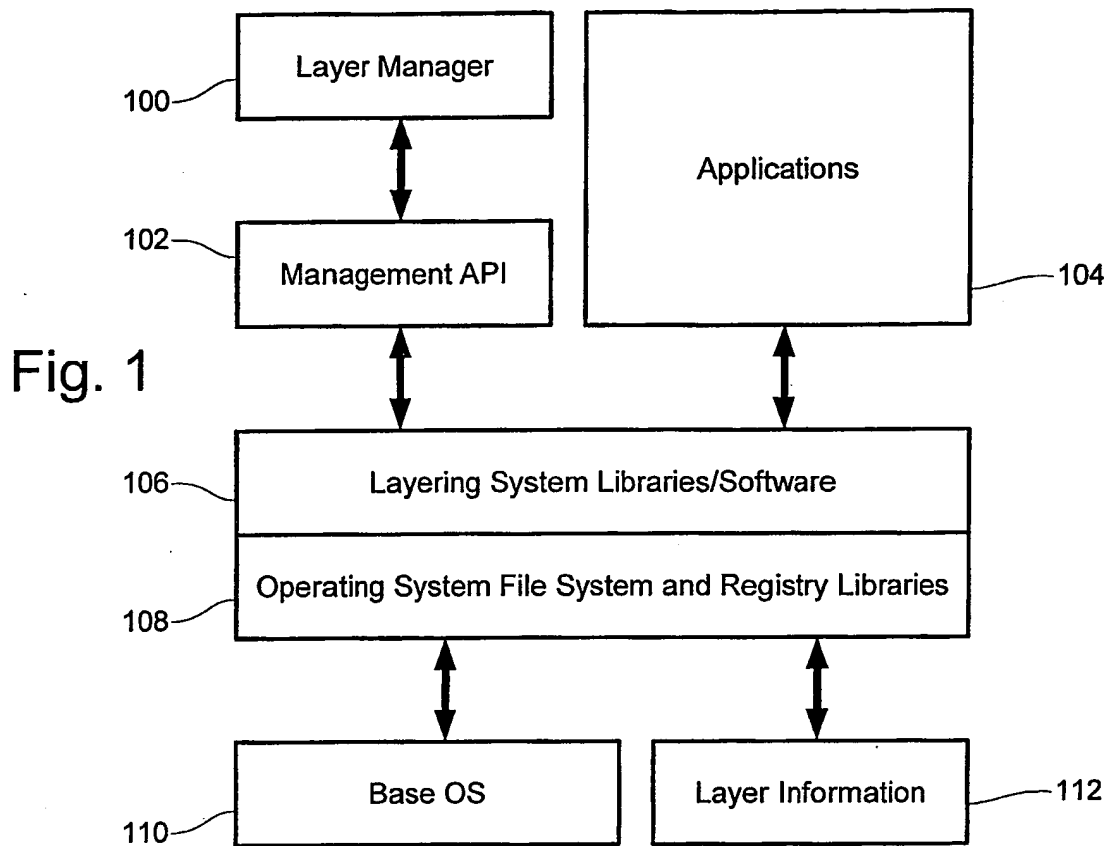



Fig. 2



2/5

Fig. 3



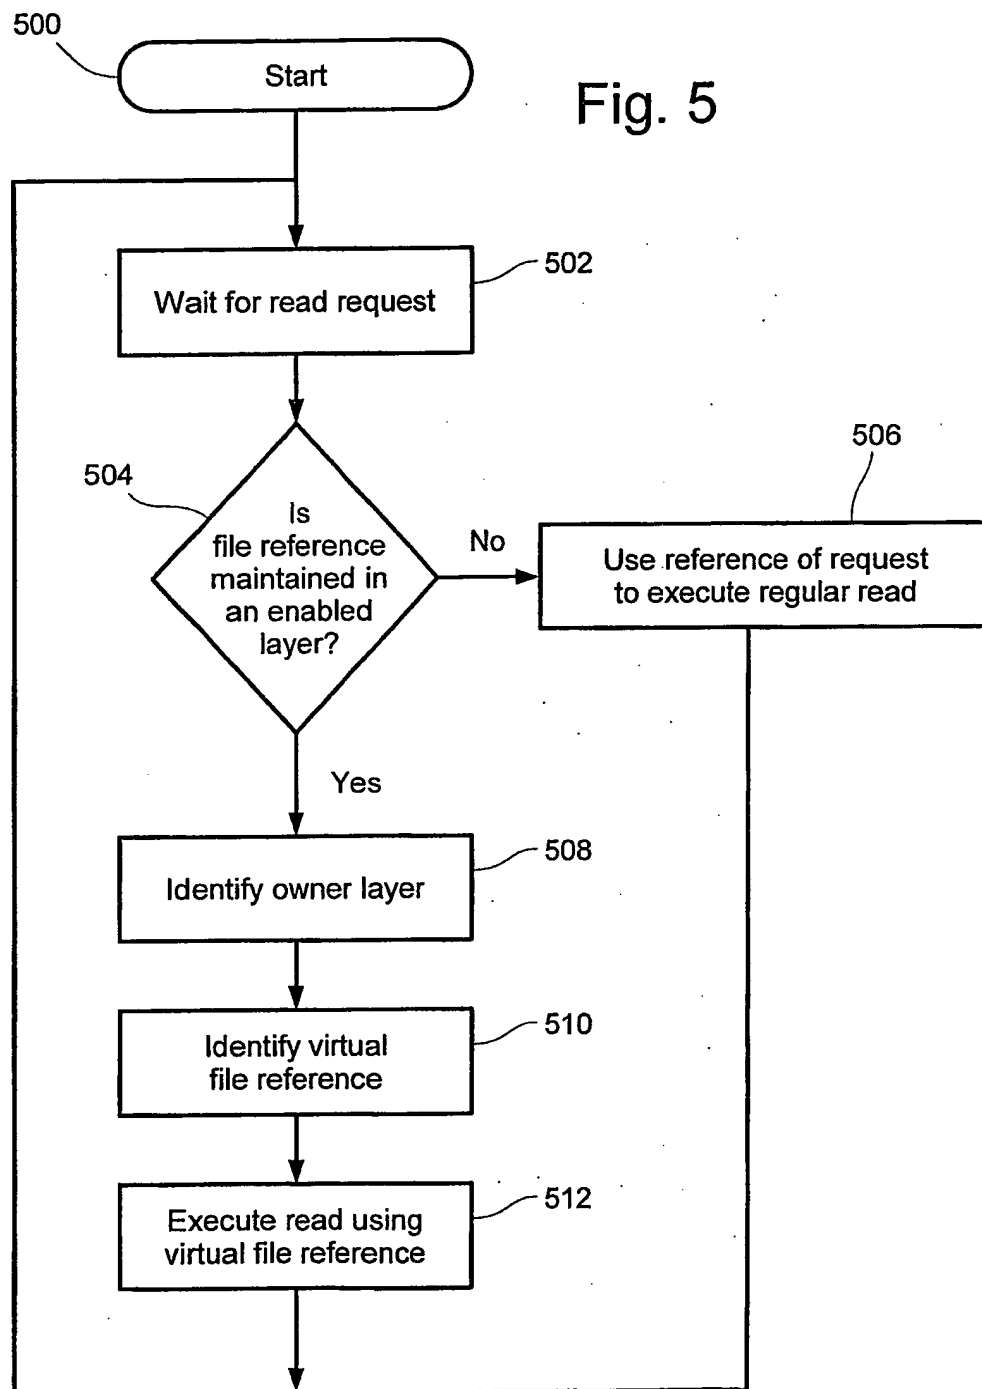


Fig. 4



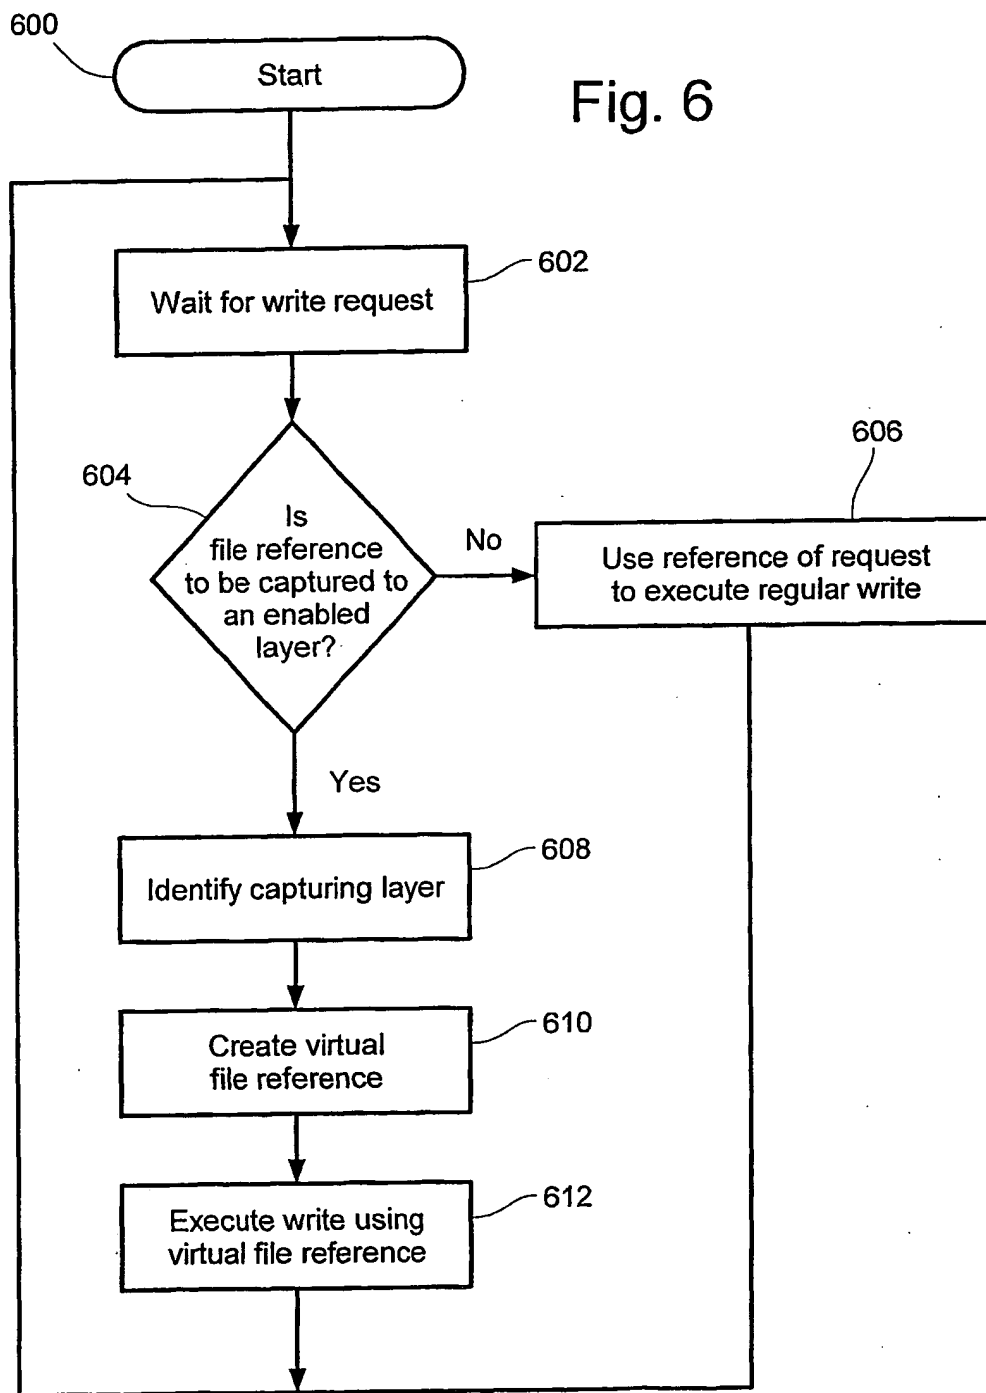
3/5

Fig. 5

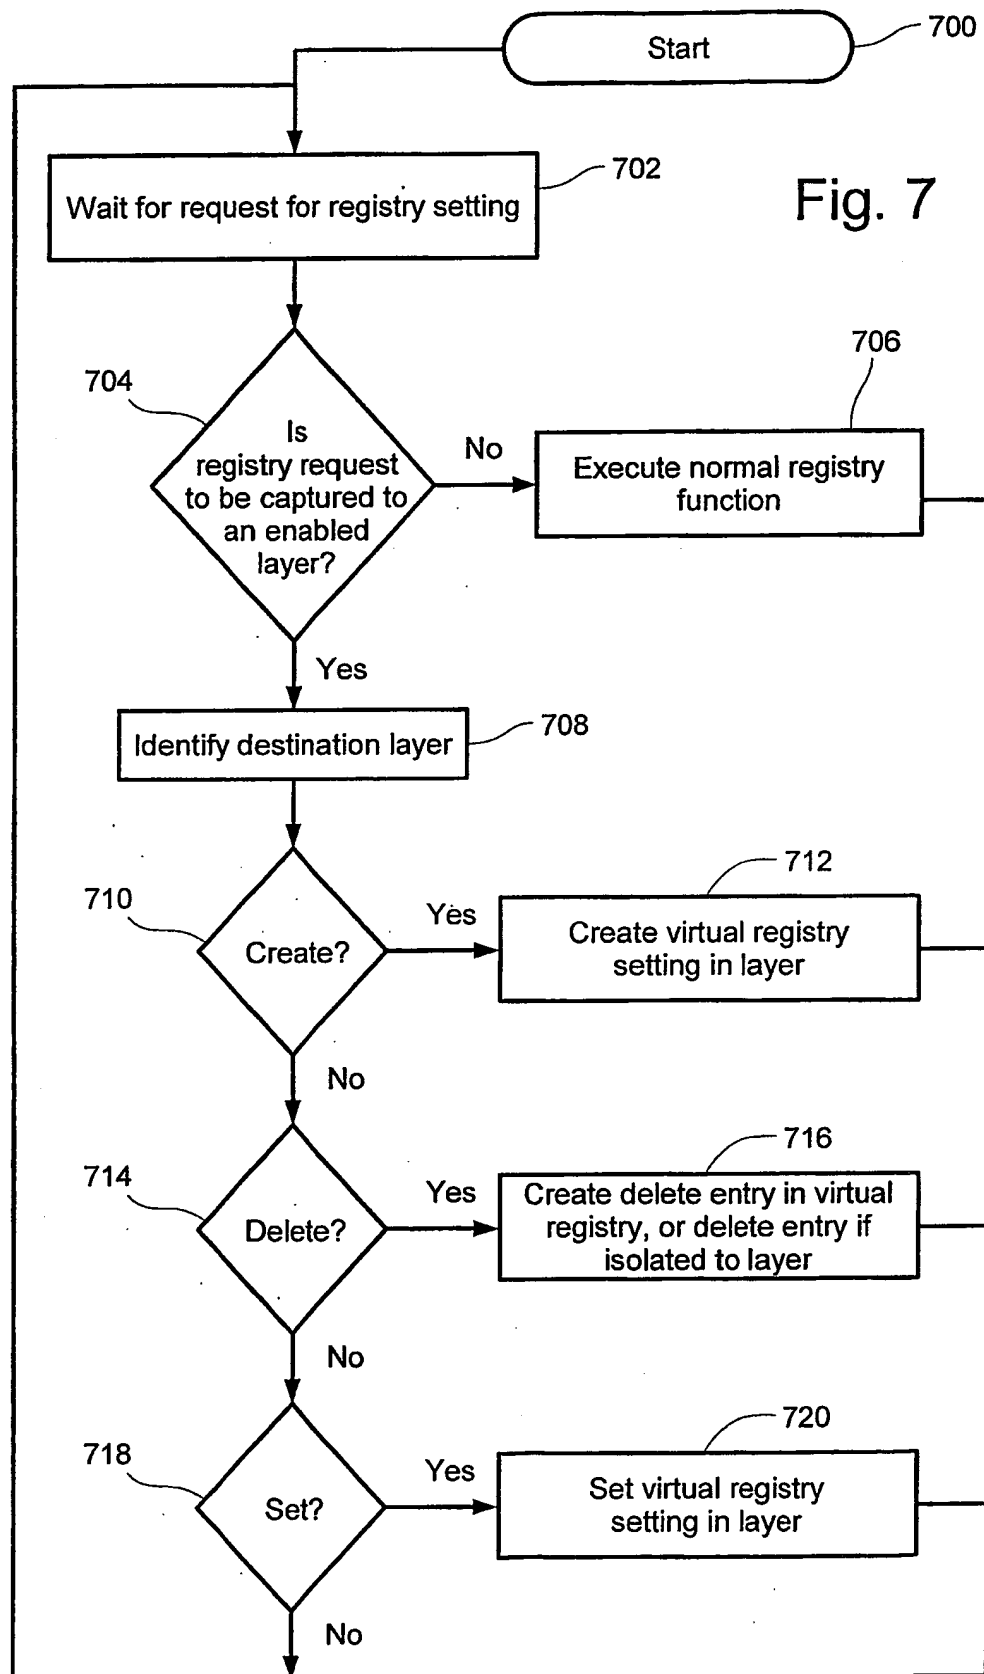


4/5

Fig. 6



5/5



INTERNATIONAL SEARCH REPORT

International application No.

PCT/US03/18633

A. CLASSIFICATION OF SUBJECT MATTER

IPC(7) : G06F 17/30
US CL : 707/200

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)
U.S. : 707/200,201,202,205

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category *	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	US 5,537,539 A (NARIHIRO) 16 July 1996, column 3, lines 20-67; column 4, lines 1-67	1-20
A	US 6,366,900 B1 (HU) 2 April 2002, column 9, lines 50-67; column 10, lines 1-16	1-20

☐ Further documents are listed in the continuation of Box C.☐ See patent family annex.

* Special categories of cited documents:

"A" document defining the general state of the art which is not considered to be of particular relevance

"E" earlier application or patent published on or after the international filing date

"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)

"O" document referring to an oral disclosure, use, exhibition or other means

"P" document published prior to the international filing date but later than the priority date claimed

"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention

"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone

"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art

"&" document member of the same patent family

Date of the actual completion of the international search

27 October 2003 (27.10.2003)

Date of mailing of the international search report

13 NOV 2003

Name and mailing address of the ISA/US

Mail Stop PCT, Attn: ISA/US
Commissioner for Patents
P.O. Box 1450
Alexandria, Virginia 22313-1450

Facsimile No. (703)305-3230

Authorized officer

Apu Mehta

Telephone No. 7036054240

(19) World Intellectual Property
Organization
International Bureau



(43) International Publication Date
24 December 2003 (24.12.2003)

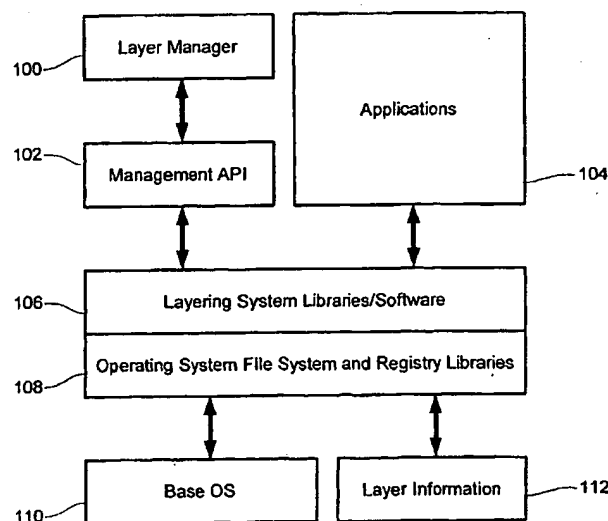
PCT

(10) International Publication Number
WO 2003/107221 A1

- (51) International Patent Classification⁷: **G06F 17/30**
- (21) International Application Number:
PCT/US2003/018633
- (22) International Filing Date: 11 June 2003 (11.06.2003)
- (25) Filing Language: English
- (26) Publication Language: English
- (30) Priority Data:
60/387,969 12 June 2002 (12.06.2002) US
- (71) Applicant (for all designated States except US): **FS-LOGIC INC.** [US/US]; 363 North University Ave., Suite 112, Provo, UT 84601 (US).
- (72) Inventors; and
- (75) Inventors/Applicants (for US only): **BLASER, Jared, R.** [US/US]; 419 North Palisade Drive, Orem, UT 84097 (US). **COOK, Randall, R.** [US/US]; 894 South 1650 East, Springville, UT 84663 (US).
- (74) Agent: **ROBINSON, Everett, D.**; Parsons Behle & Latimer, 201 South Main Street, Suite 1800, Post Office Box 45898, Salt Lake City, UT 84145-0898 (US).
- (81) Designated States (national): AE, AG, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BY, BZ, CA, CH, CN, CO, CR, CU, CZ, DE, DK, DM, DZ, EC, EE, ES, FI, GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MA, MD, MG, MK, MN, MW, MX, MZ, NO, NZ, OM, PH, PL, PT, RO, RU, SC, SD, SE, SG, SK, SL, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, YU, ZA, ZM, ZW.
- (84) Designated States (regional): ARIPO patent (GH, GM, KE, LS, MW, MZ, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HU, IE, IT, LU, MC, NL, PT, RO, SE, SI, SK, TR), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).
- Published:
— with international search report
— with amended claims and statement
- Date of publication of the amended claims and statement:
18 March 2004

[Continued on next page]

(54) Title: **LAYERED COMPUTING SYSTEMS AND METHODS**



(57) Abstract: The inventions relate generally to computer systems having facilities for providing virtual portions of file systems and configuration settings to applications. More particularly, the inventions relate to computer systems that provide a layer organization for files and configuration settings that can be overlaid on top of an operating system. An exemplary system includes a base operating system (110), operating system file system and registry libraries (108), layering system libraries and/or software (106), applications (104), a management applications programmer interface (102), and a layered manager application (100). Detailed information on various example embodiments of the inventions are provided in the Detailed Description below, and the inventions are defined by the appended claims.

WO 2003/107221 A1



For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.

AMENDED CLAIMS

[received by the International Bureau on 9 January 2004 (09.01.04);
original claims 1-20 replaced by amended claims 1-20 (8 pages)]

1. A set of computer readable media containing computer instructions for operating a layered computing environment, the set of computer readable media comprising at least one medium upon which is stored the computer instructions executable by a computing system to achieve the functions of:

(i) receiving from applications read requests for read operations to a file system, each read request containing a file reference appropriate to a file system organization,

(ii) for file references of the received read requests, attempting to identify an owner layer from a set of enabled layers;

(iii) if, for a particular read request, an owner layer is identified for a file reference, identifying a virtual read reference corresponding to the read reference of the particular read request utilizing information contained in the identified owner layer,

(iv) following the identifying a virtual read reference, causing a read operation to execute using that virtual read reference,

(v) if, for a particular read request, an owner layer is not identified for a file reference, causing the read operation to execute using the file reference of the particular read request.

2. A set of computer readable media according to claim 1, wherein the computer instructions are further executable to achieve the function of decrypting encrypted information referenced by the virtual read reference.

3. A set of computer readable media according to claim 1, wherein the computer instructions are further executable to achieve the functions of:

(vi) receiving from applications write requests for write operations to a file system, the write requests each containing a file reference,

(vii) for file references of particular received write requests, attempting to identify an owner layer from a set of currently enabled layers,

(viii) if, for a particular write request, an owner layer is identified for the file reference of that write request, identifying a virtual write reference utilizing information contained in the identified owner layer, the virtual write reference corresponding to the file reference of the particular write request,

(ix) following the identifying a virtual write reference, causing a write operation to execute using that virtual write reference,

(x) if, for a particular write reference, an owner layer is not identified for the file reference of that write request, causing a write operation to execute using the file reference of the particular write request.

4. A set of computer readable media according to claim 3, wherein the computer instructions are further executable to achieve the function of copying an existing file to the capture layer prior to causing the write operation to execute if the write operation is a modify operation.

5. A set of computer readable media according to claim 3, wherein the computer instructions are further executable to achieve the function of encrypting the information to be written prior to the execution of the write operation.

6. A set of computer readable media according to claim 3, wherein the computer instructions are further executable to achieve the functions of:

(xi) receiving from applications requests to create, delete, and set the value of a registry setting;

(xii) following receipt of a request to create, delete or set the value of a registry setting, a determining whether or not the registry operation of the request is to be captured to an enabled layer;

(xiii) acting on a request to create a registry setting, and on a determination that a registry setting is not to be captured to an enabled layer, causing the registry setting operation to execute in the base system registry location;

(xiv) acting on a request to create a registry setting, and on a determination that a registry setting is to be captured to an enabled layer, identifying a registry creation destination layer;

(xv) following said identifying a registry creation destination layer, causing the registry setting to be created virtually in the registry creation destination layer;

(xvi) acting on a request to delete a registry setting, and on a determination that a registry setting is not to be captured to an enabled layer, causing the registry setting deletion operation to execute in the base system registry location;

(xvii) acting on a request to delete a registry setting, and on a determination that a registry setting is to be captured to an enabled layer, identifying a registry deletion destination layer;

(xviii) following said identifying a registry deletion destination layer, causing the registry setting to be deleted virtually in the registry deletion destination layer;

(xix) acting on a request to set a registry setting, and on a determination that a registry setting is not to be captured to an enabled layer, causing the registry setting operation to execute in the base system registry location;

(xx) acting on a request to set a registry setting, and on a determination that a registry setting is to be captured to an enabled layer, identifying a registry setting destination layer; and

(xxi) following said identifying a registry setting destination layer, causing the registry setting to be created virtually in the registry setting destination layer.

7. A set of computer readable media according to claim 6, wherein the computer instructions are further executable to achieve the function of:

(xxii) receiving management commands through an applications programmer interface.

8. A set of computer readable media according to claim 7, wherein the computer instructions are further executable to achieve the functions of:
- (xxiii) receiving a management command through an applications programmer interface to import a specified layer archive; and
 - (xxiv) import the specified layer archive.
9. A set of computer readable media according to claim 7, wherein the computer instructions are further executable to achieve the functions of:
- (xxiii) receiving a management command through an applications programmer interface to enable a layer group; and
 - (xxiv) enable each layer of the specified layer group.
10. A set of computer readable media according to claim 7, wherein the computer instructions are further executable to achieve the functions of:
- (xxiii) receiving a management command through an applications programmer interface to disable a layer group; and
 - (xxiv) disable each layer of the specified layer group.
11. A set of computer readable media according to claim 7, wherein the computer instructions include a layer manager application.
12. A set of computer readable media according to claim 6, wherein the computer instructions are further executable to achieve the functions of:
- (xii) examining the installed layers for a configuration element, that element specifying for each layer whether or not the layer is to be enabled on system initialization; and
 - (xiii) enabling those layers having configuration elements specifying layer enablement on system initialization.

13. A computing system supporting a layered computing environment, comprising:
- a processor,
 - a storage device group accessible by said processor, said storage device group containing at least one storage device operable to contain operating system files, applications and layers;
 - one or more layers stored to said storage device group;
 - instructions stored to said storage device group, said instructions being further executable by said processor to achieve the functions of:
 - (i) receiving from applications read requests for read operations to a file system, the read requests each containing a file reference,
 - (ii) for file references of particular received read requests, attempting to identify an owner layer from a set of currently enabled layers,
 - (iii) if, for a particular read request, an owner layer is identified for the file reference of that read request, identifying a virtual read reference utilizing information contained in the identified owner layer, the virtual read reference referencing a corresponding file structure to the file reference of the particular read request,
 - (iv) following an identification of a virtual read reference, causing a read operation to execute using that virtual read reference,
 - (v) if, for a particular read request, an owner layer is not identified for the file reference of that read request, causing a read operation to execute using the file reference of the particular read request,
 - (vi) receiving from applications write requests for write operations to a file system, the write requests each containing a file reference,
 - (vii) for file references of particular received write requests, attempting to identify an owner layer from a set of currently enabled layers,
 - (viii) if, for a particular write request, an owner layer is identified for the file reference of that write request, identifying a virtual write reference utilizing information contained in the identified owner layer, the virtual write reference corresponding to the file reference of the particular write request,
 - (ix) following the identifying a virtual write reference, causing a write operation to execute using that virtual write reference,
 - (x) if, for a particular write reference, an owner layer is not identified for the file reference of that write request, causing a write operation to execute using the file reference of the particular write request.

14. A computing system according to claim 13, wherein the computer instructions are further executable to achieve the functions of:

(xi) receiving from applications requests to create, delete, and set the value of a registry setting;

(xii) following receipt of a request to create, delete or set the value of a registry setting, a determining whether or not the registry operation of the request is to be captured to an enabled layer;

(xiii) acting on a request to create a registry setting, and on a determination that a registry setting is not to be captured to an enabled layer, causing the registry setting operation to execute in the base system registry location;

(xiv) acting on a request to create a registry setting, and on a determination that a registry setting is to be captured to an enabled layer, identifying a registry creation destination layer;

(xv) following said identifying a registry creation destination layer, causing the registry setting to be created virtually in the registry creation destination layer;

(xvi) acting on a request to delete a registry setting, and on a determination that a registry setting is not to be captured to an enabled layer, causing the registry setting deletion operation to execute in the base system registry location;

(xvii) acting on a request to delete a registry setting, and on a determination that a registry setting is to be captured to an enabled layer, identifying a registry deletion destination layer;

(xviii) following said identifying a registry deletion destination layer, causing the registry setting to be deleted virtually in the registry deletion destination layer;

(xix) acting on a request to set a registry setting, and on a determination that a registry setting is not to be captured to an enabled layer, causing the registry setting operation to execute in the base system registry location;

(xx) acting on a request to set a registry setting, and on a determination that a registry setting is to be captured to an enabled layer, identifying a registry setting destination layer; and

(xxi) following said identifying a registry setting destination layer, causing the registry setting to be created virtually in the registry setting destination layer.

15. A computing system according to claim 13, wherein the computer instructions are further executable to achieve the function of maintaining a delete list.

16. A computing system according to claim 13, wherein the computer instructions are further executable to achieve the function of executing a lazy thread.

17. A computing system according to claim 13, wherein the computer instructions are further executable to achieve the function of decrypting encrypted information referenced by a virtual read reference.
18. A computing system according to claim 13, wherein:
the computer instructions are further executable to achieve the function of (xi) receiving management commands through an applications programmer interface; and
the computer instructions include a layer manager application.
19. A set of computer readable media according to claim 18, wherein the computer instructions are further executable to achieve the functions of:
(xii) receiving a management command through an applications programmer interface to enable or disable a layer group; and
(xiii) enable or disable each layer of the specified layer group.

20. A method of handling requests to read and write files in a layered computing system, the method comprising the steps of:

- (i) receiving from applications read requests for read operations to a file system, the read requests each containing a file reference,
- (ii) for file references of particular received read requests, attempting to identify an owner layer from a set of currently enabled layers,
- (iii) if, for a particular read request, an owner layer is identified for the file reference of that read request, identifying a virtual read reference utilizing information contained in the identified owner layer, the virtual read reference referencing a corresponding file structure to the file reference of the particular read request,
- (iv) following an identification of a virtual read reference, causing a read operation to execute using that virtual read reference,
- (v) if, for a particular read request, an owner layer is not identified for the file reference of that read request, causing a read operation to execute using the file reference of the particular read request,
- (vi) receiving from applications write requests for write operations to a file system, the write requests each containing a file reference,
- (vii) for file references of particular received write requests, attempting to identify an owner layer from a set of currently enabled layers,
- (viii) if, for a particular write request, an owner layer is identified for the file reference of that write request, identifying a virtual write reference utilizing information contained in the identified owner layer, the virtual write reference corresponding to the file reference of the particular write request,
- (ix) following the identifying a virtual write reference, causing a write operation to execute using that virtual write reference,
- (x) if, for a particular write reference, an owner layer is not identified for the file reference of that write request, causing a write operation to execute using the file reference of the particular write request.

Statement under Article 19(1)

Applicant, by these newly presented claims, seeks only to clarify the claimed subject matter, and does not amend these claims for any reasons of patentability. Applicant believes both the originally filed claims and these newly presented claims are patentable over all known prior art.